


RESEARCH ARTICLE

The 2020s political economy of machine translation

Steven Weber^{1*}  and Nikita Mehandru¹

¹University of California Berkeley, School of Information, Berkeley, CA 94720, USA.

*Corresponding author: Steven Weber, email: steve_weber@berkeley.edu

Abstract

This article specifies and explores the hypothesis that the diversity of human languages, right now a barrier to “interoperability” in communication and trade, will become significantly less of a barrier as machine translation technologies are deployed over the next several years. We argue that machine translation will become the 2020’s analogy for ideas, to what container shipping did for goods trade in the second half of the twentieth century. But as with container shipping or railroads in the nineteenth century, this new boundary-cost and transaction-cost reducing technology does not reduce all boundary and transaction costs equally, and so creates new challenges for the distribution of ideas and thus for innovation and economic growth. How we develop, license, commercialize, and deploy machine translation will be a critical determinant of its impact on trade, political coalitions, diversity of thought and culture, and the distribution of wealth.

Keywords: trade; globalization; machine translation; inequality; productivity

The most important and most desirable technology revolutions are those that boost economic productivity and improve the human experience at the same time. That normative assertion sounds obvious, but historically it sets a high bar. The eighteenth-century Industrial Revolution would not qualify because for most people outside of factory-owning capitalists, productivity rose at the cost of a largely degraded human experience in many aspects of life. The introduction of electricity probably comes closer to qualifying, as does the printing press. There’s long been doubt about whether digital technologies, computing, and the internet meet these criteria. The productivity gains are unmistakable, but the overall impact on human experience is now generally seen as ambiguous and possibly net negative.¹ How we understand, organize, deploy, pay for, and incorporate technology is what determines its impact. Those elements are largely under the control of people but are much easier to maneuver at early stages rather than later in the life cycle of any particular technology.

The core proposition of this article is that machine translation is just now entering the early “sweet spot” period where thoughtful decisions can tilt the table toward better impact for both economic productivity and human experience. A subsidiary proposition is that because machine translation runs on a global computing infrastructure, it is inherently necessary to think from the start about its global impact (the steam engine took decades to diffuse from its British home to the rest of the world).

The broadest ambition should be to develop and deploy machine translation technology in a manner that is Pareto-improving across five dimensions of liberal progress: enhancing productive competition, ensuring fundamental fairness, guaranteeing a baseline of protection for the most vulnerable,

This work was supported in part by the Hewlett Foundation under a grant to the Center for Long Term Cybersecurity at Berkeley and by the Berggruen Institute (Los Angeles) where the author is a Berggruen Fellow.

¹Boulianne, S. and Y. Theocharis. (2020), 111–27.

© The Author(s), 2021. Published by Cambridge University Press on behalf of V.K. Aggarwal

deepening pluralism, and stimulating innovation.² A narrower goal is to aim simply for greater equality in the distribution of economic gains, as a necessary but not sufficient condition.

Put differently, the argument is machine translation will not contribute to advancing global liberal outcomes if it reinforces (or supercharges) contemporary trends toward imbalanced distribution of economic gains from digital technology. But in the absence of intentional intervention, We argue that this is exactly what is set to happen. We can do better, and the normative goal of this research is to reason toward specific proposals that would improve both productivity and distributional equity at the same time.

Language matters

Human language is central in the coordination of complex economic activities that generate productivity. Consider a (possibly fictitious) example from the biblical story of the Tower of Babel:

now the whole world had one language and a common speech ... they said, come let us build ourselves a city with a tower that reaches to the heavens that we may make a name for ourselves.

Long before economics had concepts of transaction and coordination costs, the biblical story had a deep intuition about the importance of language, written and spoken, in both. The bible overestimated what bricks and mortar could do, but it does recognize the potent combination of human ambition and the ability to coordinate. The story might be seen as an early version of the modern dictum “if we really work together there is nothing we cannot do.” In this case, having a single language would enable coordinating the labor of all humankind so that people could build a tower to heaven and come close to being on the same level as God.

God didn’t much like that idea. And so:

the Lord came down to see the city and the tower the people were building. The Lord said, if as one people speaking the same language they have begun to do this, then nothing they plan to do will be impossible for them. Come, let us go down and confuse their language so they will not understand each other.

Thus, the introduction of language-based transaction costs friction into economic coordination, which “scattered the people over the face of the whole earth.” If groups of people couldn’t communicate with others to coordinate production, there wasn’t much reason for them to stay in the same physical location. People who spoke the same language went off in groups to inhabit different parts of the earth; production of towers and other things reverted to smaller scale coordination challenges; and distinctive cultures developed in somewhat insulated environments where ideas from “inside” a particular language could diffuse orders of magnitude more quickly and easily than ideas from an “outside” language silo that was both physically and linguistically separate.³

The Babel story highlights a few important arguments—upside and downside—about the multiplicity of human languages. First, even moderately complex economic coordination requires granular communication. Introducing language friction into that process is a significant transaction cost in itself, and probably magnifies the impact of other transaction costs at the same time (e.g., costs associated with imperfect contracting likely are made worse when contracts are rendered in multiple languages). Second, language barriers have a protective effect on the evolution of different ideas, practices, religions, and culture, that might otherwise be subject to homogenization or at least consolidation if those barriers disappeared (more on that later).

²These five dimensions are drawn from forthcoming work with Naazneen Barma and Brent Durbin, *How to Make the World More Liberal*. Oxford University Press, forthcoming 2022.

³The analogy here is to “evolutionary isolation” as explained by Charles Darwin’s study of the Galapagos Islands. Mayr (1959), 221–30.

Third, the impact on overall productivity of a multiplicity of languages is at the highest level a consequence of how these two vectors intersect. More diversity of ideas should contribute positively to innovative potential and so would help increase productivity (multiple experiments going on in parallel accelerating useful discovery); but the language barrier to coordination becomes a drag on learning from experiments in other places where what is discovered gets encoded in a different language. It also constrains the ability to scale.

There's no *a priori* argument to determine which vector dominates over time, but intuition along with a casual reading of history suggests that learning across cultures and languages is much slower than it "ought" to be; that mistakes are repeated and dead-ends pursued much more often than they should be; and that increased diversity of ideas and practices relevant to economic productivity isn't sufficient to overcome those negatives. If that intuition is roughly accurate, then overall economic productivity grew more slowly because of language diversity (which is, after all, what God intended in the biblical story). That's not a comment on whether cultural and other forms of diversity are "worth it," it's simply a hypothesis that productivity would have been higher in an alternative history where—all other things equal—human beings spoke the same language.

The two hypotheses about the conflict-promoting or conflict-avoiding effects of translation map almost exactly onto long-standing arguments in international relations theory about interdependence between countries. Commercial liberalism posits that higher levels of trade between countries bind them together in valuable relationships that conflict would place at risk, and thus interdependence makes conflict less likely. Structural realism, in contrast, argues that high levels of interdependence are a source of conflict (over distribution of gains, technology appropriation, and so on) and that lower levels of interdependence are more likely to be associated with peace because countries that are less dependent on each other, have less opportunity to try to use the resulting leverage against each other and have less to fight about overall.⁴

Rather than one theory being right and the other wrong, what is more accurate is that one or the other causal mechanism predominates under different sets of conditions. It's reasonable to posit that the same is likely to be true of translation technologies—that the overall effect of translation on conflict (as one example of an outcome) will be some kind of sum of multiple causal mechanisms and thus depend on conditions that amplify or weaken upside and downside vectors.

That's a hopeful inference because some of those conditions are going to be subject to decisions that people make about how the technology is developed, deployed, licensed, paid for, and so forth. But before we can get more specific about what those conditions and decisions look like, we need to examine some of the constraining affordances under which modern machine translation technology operates.⁵ How the technology works and what those mechanisms do and do not make possible in the medium-term future set the stage for what decisions can be realistically framed up.

“Universal” isn't universal

Ethnologue currently estimates that there are more than 7,000 known living languages, many of which are used by very small numbers of people. In comparison, Google's Cloud Natural Language API recognizes less than 100 languages and, as of February 2020, offers syntactic analysis on 11 and sentiment analysis on 10.⁶

The perennial goal of universal machine translation is to create a single model that translates between any and all language dyads with equally high accuracy. Today's technology isn't close. A casual user can see some of the limitations by playing with Google Translate on any web browser. Google Translate offers about 100 languages; it works extremely well translating between English

⁴See, e.g., Doyle (2016; The microfoundations of these effects have been hypothesized in more elaborate ways over time (such as multinational firms who depend on global supply chains lobbying governments against paths to conflict).

⁵A relevant analogy around affordances can be found in Adler-Nissen, Drieschova (2019), 531–45.

⁶Language Support [User Documentation]. <https://cloud.google.com/natural-language/docs/languages> (accessed August 26, 2021).

and other common languages; moderately well between English and less commonly used languages; and less well between a pair of infrequently used languages. Some of the spoken language translation services (e.g., now available on mobile phones) are even more limited. These systems seem almost magical if you are an English-only speaker working with French in Paris or Spanish in Madrid; they are much less useful if you are a Burmese-only speaker trying to communicate in Basque.

The past decade has seen continuous improvement in machine translation techniques. Unlike rule-based and statistical machine translation methods, neural machine translation (NMT) does not need human-designed features. NMT provides an end-to-end framework where the model can take a source sentence and generate the next target sentence word by word.⁷ Some of the largest technology companies have begun to tackle very tough machine translation problems: Microsoft was able to achieve human parity on Chinese-to-English news translation; Facebook and Google have heavily invested in translating some of the world's lowest-resourced languages. Because NMT techniques are more sensitive to data quality though, compared to statistical and phrase-based machine translations, there are significant challenges in ensuring equally distributive economic gains.⁸

The simplest way to understand the limitations of conventional systems is to think of today's machine translation algorithms as supervised learning systems that are built on reliable and accurate labeled training datasets. As is often true, high-quality training datasets are out in the world as a coincidence of other needs and processes. Commonly used datasets for machine translation include TED talks (manually translated in up to 59 languages); documents from the European Parliament (which by law are translated into 21 European languages); and the UNCorpus (more than 11 million sentences rendered in six languages). The Bible Corpus is generally considered the broadest public available multilingual dataset (about 30,000 sentences in 900 languages).⁹

Linguists and machine translation scientists make an important distinction between what they call high-resource and low-resource languages.¹⁰ High-resource languages are, expectedly, languages for which many data resources exist—English is by the far the highest resource language, with many Western European languages as well as Japanese and Chinese also reasonably high resource. Low-resource languages—including many languages from poorer countries, many local dialects, rarely used or semi extinct languages—make up the vast majority of the rest. There are a large number of languages that are mostly spoken and for which very few written resources exist. There are languages with a large amount of digital raw text from various genres (everything from social media to scientific papers) and lexical, syntactic, and semantic resources (dictionaries, semantic databases) as well as various bodies of highly annotated text (e.g., text that is labeled with parts-of-speech tags); and many languages that have little to no such resources available.

This resource gradient among languages and the resulting disparity in machine translation capabilities is the most important technology variable shaping how machine translation will transform economies and culture. State-of-the-art NMT systems aspire to develop a single multilingual model that could handle all translation directions at once, rather than relying on a large number of dyadic models that could on aggregate, at least in theory, do the same (the number of necessary dyads would scale as a square of the number of languages; that many models would be almost impossible to train, deploy, and maintain). Translating in sequential dyads (e.g., going from Swahili to English and then from English to Spanish, to translate from Swahili to Spanish) is computationally intensive and less accurate, as some level of error is introduced at each sequential step.¹¹

⁷Wang et al (2021).

⁸Awadalla et al. (2018).

⁹Cettolo, Girardi, & Federico (2012), 261–68; Koehn (2005), 79–86; Ziemski, Junczyz-Dowmunt, & Pouliquen (2016), 3530–34; Tiedemann, J. (2018), 188–97.

¹⁰See, e.g., Maimaiti et al. (2019).

¹¹Background explanations in Wu, Y., M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, L. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. S. Corrado, M. Hughes, and J. Dean. 2016. "Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation." ArXiv, abs/1609.08144.

A 2019 paper from Google AI surveys the state of universal machine translation.¹² What would make the problem easier would be if the learning signal from any particular language had a positive impact on the ability of the model to work with other languages—put differently, if the model were to generalize more effectively with each language that it learns. This is called positive transfer and it does appear to some degree in low-resource languages; but the gains start to reverse themselves and performance on high-resource languages start to decline after a certain point (this is called negative transfer or interference).

The massively ambitious model described in that paper uses an original dataset extracting parallel sentences from the web and contains 25 billion sentence pairs; and dataset training example sizes that range from around 2 billion for high-resource language pairs down to 35,000 for low-resource language pairs, a difference of almost five orders of magnitude. The researchers experiment with several methods aimed at enhancing positive transfer and reducing interference (e.g., oversampling low-resource language dyads to compensate for the imbalance in the training data). But there is no free lunch—oversampling on low-resource languages improves transfer but creates interference that reduces performance significantly on high-resource languages; regular sampling yields better retention of performance on high-resource languages but sacrifices considerable performance on low-resource languages. More sophisticated sampling strategies change the terms of these trade-offs but do not eliminate them. The same is true for increases in model capacity that rely on better and more hardware and infrastructure.

What this adds up to is a technology weighted toward high-resource language dyads and against low-resource language dyads, with high-low dyads achieving translation capacity somewhere in the middle. Barring an unspecified technology breakthrough that would change the basic methods and remake the terms of the trade-offs, it's reasonable to extrapolate substantial improvements in all translation tasks over the next decade, *but not all equally*. This imbalance—possibly increasing in relative magnitude—between translation among high-resourced languages as compared to other dyads, will become a crucial feature of how this technology reshapes communication, commerce, and culture on the global landscape.

Historical analogies and interoperability

Translation today is more like interoperability than it is integration or uniformity of language. Integration implies a frictionless state; uniformity implies a single language (like Esperanto). The IEEE defines interoperability as “the ability of two or more systems or components to *exchange* information and to *use* the information that has been exchanged.”¹³ The key distinctions are nicely clear when it comes to computer languages. Technical interoperability is the ability to open a file created in one application within a second application. Syntactic interoperability is a common format for data exchange (XML is an example). Semantic interoperability is the ability to understand the meaning of content in the same way by sender and receiver. Pragmatic interoperability is the ability to do things (e.g., coordination) that follows from semantic interoperability and would not be possible without it.

These distinctions matter because they characterize the kinds of frictions that machine translation's partial interoperability capacity will reduce, leave in place, and exacerbate or create anew. Not all aspects of interoperability depend on each other nor do they reinforce each other in linear ways. A casual but illuminating example that many people have experienced comes from the ability to navigate a subway map in a foreign country. In a recent trip to Japan with a partner who speaks a small amount of Japanese (we speaks none), one of us was able to plan the subway trip around Tokyo much more easily than the partner with some language skills. That is because it was possible to navigate the map (which was labeled only in Japanese) through syntax, interpreting the “common form” of a subway map without translating a single label. The partial Japanese speaker had the capacity for partial

¹²Arivazhagan et al. (2019), 1–27.

¹³Interoperability. 1990. In *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*. IEEE. Italics added.

semantic interoperability, which led her to focus on the labels and *reduced* her pragmatic interoperability (the non-Japanese speaker's capacity was sufficient to get where we needed to go, without any semantic interoperability at all).

The same strategy might not work with pure text, but because machine translation will be applied to all kinds of artifacts that have words on them as well as pictures and diagrams (think of a technical manual for a machine tool) the story is still instructive. At the limit, the various kinds of interoperability might be synergistic with each other; but short of the limit (and where we will be in the next decade), it often won't be so.

This argument contextualizes the relevance of three simple historical analogies. The use of analogies here is a heuristic for starting to characterize challenges that machine translation will pose to global political economy. The first analogy is to what in the 1990s was called globalization theory. In simple terms, the proposition is that language may now be as or more significant than physical distance as a barrier to interoperability in communication and trade across borders (We refine and evaluate that heuristic later).

The second analogy is about previous generations of technology that reduced the impact of those physical barriers in selective ways—notably railroads in the nineteenth century and container shipping in the twentieth century. Both had hub-and-spoke topologies that shaped their impact on trade and as a consequence economic growth. This phenomenon is well documented in nineteenth-century railroad routes and pricing schemes. It was, for example, very cheap to ship goods on rail from Chicago to St. Louis; more expensive to ship from St. Louis to non-hub small cities (like Dubuque, Iowa); and very expensive to ship between two non-hub small cities (from Dubuque to Rock Island, Illinois). This was partly determined by the technology of rail and partly by the business model of railroads, but the consequences were equivalent (at least until antitrust authorities forced reform of what were determined to be discriminatory pricing schemes).¹⁴

Railroads vastly reduced the impact of physical distance *between particular nodes* in a network and, in relative terms at least, *increased* the barriers between others. This had the consequence among other things of further concentrating economic activity and exchange in the large city nodes and hollowing out of small towns and communities along the way.

The third analogy is about *The Mythical Man-Month* and what later became known as Brooks's Law (after the book's author Frederick W. Brooks).¹⁵ Brooks was an observant software engineer at IBM in the 1970s who sought to explain why it seemed so difficult for teams to coordinate and cooperate effectively in the creation of complex software products. His view was distilled into Brooks's Law: that adding additional person-power to a late software project only makes it later. Brooks posited several mechanisms to explain this observation. The most obvious was the ramp-up time for a new engineer to understand the work done before she joined the team (and the effort that team members would need to expend to help her get there).

More relevant to the machine translation challenge was Brooks's argument that communication overhead (a form of transaction costs) increases at a very rapid rate when the number of people working together on a complex task rises. As Brooks put it, as more people are added to a team each person has to spend more time figuring out what everyone else is doing. That dynamic is less problematic for tasks that are easily divisible into discrete parts (which is, of course, Adam Smith's argument for a classical division of labor). But when tasks are highly interdependent or need to be precisely sequenced or really can't be divided up into discrete elements and then reassembled, Brooks's Law says something profound about the vagaries of human communication and why the division of labor is not always a boost for productivity. As Brooks put it so cleverly, the fact that one person can produce a baby in nine months does not mean that nine people working together can produce a baby in one month.¹⁶

¹⁴Berk (1997).

¹⁵Brooks (1975).

¹⁶The same dynamic operates in the opposite direction: The mythology says that Helen of Troy had such beauty that she could launch a thousand ships, but a milli-Helen would almost certainly not be the exact amount of beauty needed to launch one ship.

Software engineering has progressed dramatically in both technical and sociological terms since Brooks wrote. Agile development processes and code repositories and tools like Slack and many of other innovations have softened the terms of Brooks's Law, but not eliminated the underlying insight about how hard it is for humans to communicate in words about abstract concepts and artifacts like software that exist almost on another plane, and at a level of complexity that no single person can visualize in their mind at once.¹⁷

Machine translation will enter the picture with globalization dynamics, economic geography and Brooks Law still in place; the next section of the article explores how these analogies and insights illuminate key challenges for political economy on the global stage.

Challenges to global political economy

Machine translation at present poses at least three specific challenges to liberal progress in global political economy. The three we consider here are comprehension and false positives; diversity of thought and the power law; and technical barriers versus cultural barriers including elite differentiation within imagined communities.

Comprehension and false positives

Shared language isn't equivalent to shared understanding, just as technical interoperability might not be semantic or pragmatic interoperability. Intuition and experience tell us that it is remarkably easy to misinterpret the meaning behind another person's speech even when they speak the same language. Most people have had decades to calibrate themselves in this respect and we rely on social signals and nonverbal cues to "test" our understandings of what others have said in conversation. Even when we have a shared foundation of language, idioms, and compatible cultural references, it's not easy and it's far from perfect—but we get it done with reasonable efficacy.

The first decade(s) of widespread machine translation will confront an environment where fewer or (sometimes) none of these shared foundations are present. The most profound human challenge of machine translation might very well be that people using the technology will *think* they understand more than they do about each other and about what the other has said. The fact that everyone will know this is possible, or at least could know it, doesn't negate the likelihood that false positives will happen at a much higher rate.

A good analogy here is the experience in health care when MRI scans first became available on a widespread basis.¹⁸ A kind of natural experiment, it was a test of skilled peoples' ability to quickly incorporate awareness that there was no established baseline of what was "normal" anatomic variation, or at least an observation not worth doing anything about, when suddenly the resolution at which one could see inside body structures jumped. In practice, false positives became common, as a very large number of MRI scans showed what looked like abnormalities against a now obsolete baseline (generally the much less detailed X-ray). What was later understood to be normal variation, was at the beginning overinterpreted as pathology—in other words, the signal was read as having more meaning than it contained. There were costly consequences—unnecessary procedures and surgeries that exposed patients to risks without corresponding benefits.

It's easy to foresee the same dynamic unfolding in the early years of widespread machine translation usage, and particularly in spoken language where peoples' concentration levels vary widely. I will think I understand what you are saying better than I really do, perhaps by taking your words too literally or filtering them through my cultural metaphors (or weakly understood versions of yours). My false

¹⁷Meso, Jain. (2006), 19–30.

¹⁸An example of this overdiagnosis problem is Welch et al. (2010). I lived through this as a medical student and it was memorable and enlightening how many "incidentalomas" early use of MRI uncovered. Of course, the inability to parse false positives wasn't just cognitive, it was also driven by codes of practice, insurance, and potential malpractice liability concerns that took their own time catching up to the "new" knowledge environment.

positives might sometimes be curiosities only, but they could matter quite a great deal, for example, if I create expectations about what you are going to do next based on my (faulty) understanding—and then you do things that defy those expectations in ways that hurt my interests. The “right” response to that scenario would be for me to re-correlate my confidence levels in the output of the translation algorithm from the start. The likely response (more likely in cases in which you are an adversary or competitor) will be for me to infer that I’ve been deceived or manipulated by you, leading to a decline in ambient trust and higher levels of conflict.

Over time, false positives will become less frequent as baseline understandings evolve in the presence of experience (exactly what happened with MRIs). But that process could easily take a decade or more, and it’s the costs and instabilities associated with the transition period that will stand out.

Diversity of thought and The Long Tail

In 2006, Chris Anderson’s *The Long Tail* popularized the argument that the internet’s ability to reduce search costs and (for digital goods at least) the costs of holding nearly infinite inventory would lead to a flourishing of latent human tastes for an extremely broad diversity of products and ideas.¹⁹ The iconic example was supposed to be music and books. Instead of a small number of superstar titles dominating markets, the internet would create a “long tail” where small numbers of people had strong preferences for much greater diversity. The argument made intuitive sense and it slotted in nicely to idealistic visions of digital liberalism that were still common in the that decade.

The long tail effect has certainly persisted under some conditions. Brynjolfsson et al. (2011) analyzed data collected from a multichannel retailer and empirically found that even with the same products and prices, the internet channel experienced a significantly less concentrated sales distribution compared to the traditional channel. Consumers on the internet channel leveraged search and discovery tools, such as recommendation engines, which was associated with an increase in the share of niche products.²⁰ Holtz et al. (2020) found similar patterns of niche selection in their large-scale, randomized experiment on Spotify that examined the effect of personalized recommendations on consumption diversity. They found that the group given personalized podcast recommendations based on their music listening history, as opposed to the most popular podcasts among their demographic, increased the average number of podcast streams per user but decreased the average individual-level diversity of podcast streams. Overall, however, this group did experience an increase in aggregate diversity of podcast streams, signaling some long tail effect not within users but across them.²¹ Interestingly, Elberse and Oberholzer-Gee (2007) studied both the long-tail and superstar phenomena in the context of US home video consumption from 2000 to 2005 but instead found that it’s ultimately quite difficult for content providers to benefit from the long-tail effect. Specifically, the authors note that increased variety can fragment markets—less well-known titles reached smaller audiences—but that these audiences are also less loyal and when a “superstar” does come along will tend to abandon their niche.²²

Thus, the long tail phenomenon turned out to be mostly wrong. Breaking down the argument into its component parts helps to clarify important lessons that apply to machine translation. The fact that digital environments can support more variety than traditional physical markets, does not necessarily mean that they will do so. Consider the logic of the demand side. In theory, using Google instead of print or broadcast media to learn about rare products or ideas should enable much more personalized preference matching. But note the underlying assumption that there exist “original” preferences that are in fact much more varied than physical search and distribution mechanisms were able to satisfy. Similarly in theory, sophisticated personalization algorithms that help people search (advanced versions of “people who like X also like Y”) should enhance the long tail effect. But note the underlying

¹⁹Anderson (2006) *The Long Tail*; Brynjolfsson et al. (2010). “The Longer Tail: The Changing Shape of Amazon’s Sales Distribution Curve.”

²⁰Brynjolfsson et al. (2011), 1373–86.

²¹Holtz et al. (2020), 75–76.

²²Elberse (2007), 49–72.

assumption here, that personalization algorithms are tuned to do that, rather than tuned to the contrary assumption that people like X and Y not because of some unique personal taste but because they want to consume ideas that can contribute to sociability with others who have consumed similar ideas.

An alternative view is that baseline demand is *less* diverse than we might imagine, simply because most ideas are valued in a social context where people want to talk with each other. People might want to consume “superstar content” simply because that act supports social interaction. They also might use popularity as a proxy for otherwise hard-to-measure quality.²³ Combine those two elements of the demand function and it’s easy to see why superstars still dominate over the long tail and, in many markets, have increased their dominance.

The logic of the supply side might point in the same direction. What are the incentives to *create* a diversity of ideas? Technology certainly reduces the cost of “stocking” and “distributing” niche ideas, effectively to zero. But *creating* new ideas that matter is still hard and expensive. The expected return on investment in a niche idea remains very low unless and until it becomes popular.

Ben Thompson’s aggregation theory explains how the current structure of major digital markets enhances this logic.²⁴ The decline in power of gatekeepers such as TV networks and traditional publishers that once controlled access to means for the promotion and distribution of ideas didn’t create a level playing field that simply empowered individuals. Instead, it gave rise to aggregators who gained market power by pulling together and organizing demand, and then presenting that demand to suppliers. Search engines and social networks are the iconic examples.

How would these dynamics change with machine translation? It’s hard to imagine a mechanism by which reduction of language barriers within markets for ideas would reestablish power for supply-side gatekeepers. Rather, it seems likely to reinforce the power of the demand-side aggregators, who now would be in a position to aggregate demand even more effectively from a much larger set of potential consumers regardless of their native language. To the extent that ideas seek large markets as products do, then the returns to a popular idea in a larger market with fewer language barriers will be even greater.²⁵ Demand-side aggregators will be the ones sending signals to suppliers about where those concentrations lie. Their power would be enhanced even more by the fact that some of the most important aggregators—including Google—are also the leaders in machine translation technologies. And so, their investment and deployment strategies for machine translation could be tweaked in ways that enhance their power as aggregators even further. There also exists a clear incentive problem for platform providers to develop services for low-resource language speakers. Not only is the return on such an investment unclear, but the lack of data available to train machine learning models adds more friction to the process. Wills et al. (2019) approximated the commercial value of a language by multiplying the number of language speakers by the national GDP per capita of those speakers and found that the top 100 languages covered almost 96 percent of global GDP.²⁶ These 100 languages, however, account for less than 60 percent of the total world population living on less than \$1.90 per day. Thus, the languages spoken in the poorest regions of the world unsurprisingly hold little commercial value and have limited amounts of available data. In a 2008 study, Choudhury demonstrated that the resource distribution across languages follow power-law, and that four languages—Arabic, Chinese, English, and Spanish—have the largest amount of data resources while 90 percent of the world’s languages in the long tail of this distribution had little, if any, resources to train a natural language processing (NLP) system.²⁷

²³There is a significant individual cognitive burden associated with assessing new ideas; this limitation on the demand for diversity would operate independently of the desire for sociability around closely related ideas.

²⁴Thompson, B. (2015).

²⁵“Returns” could obviously be economic, or otherwise (“fame” and notoriety is not only and always a purely economic concept/value, of course).

²⁶Adam Wills and Georgia Barrie of Farm.ink and Jake Kendall of the DFS Lab

²⁷See Microsoft Research, ELLORA: Enabling Low Resources Languages Project, described at: <https://www.microsoft.com/en-us/research/project/ellora/> (accessed August 26, 2021).

The boldest hypothesis about how this would look at scale has three elements. First, the superstar phenomenon would converge into a small number of highly populated “echo chambers” of ideas, reducing overall diversity and increasing the intensity of competition (and possibly conflict) among them. It would become harder for creative new coalitions to emerge because it will be difficult for idea entrepreneurs to mobilize small groups with partially overlapping ideas and pull them together through compromise. While Axelrod (1997) found that local convergence can sometimes preserve global diversity, the superstar phenomenon toward high-resource languages would instead occur and be accelerated by the fact that there tends to be a premium on learning such languages.²⁸ English, for example, has the value of *lingua franca*, particularly in previously colonized countries, and gaining fluency can serve as a signaling effect. Second, there would be enhanced incentives to create ideas that have the potential to become globally popular—but vanishingly few would be able to do so and the vast majority would get little traction and likely die. Third, the strongest incentives will probably be to create “complementary ideas” (as well as products) that naturally appeal to already existing superstar idea clusters (because these are the biggest preexisting markets primed for complementary consumption).

Another way to think about that argument is to highlight the attractiveness to idea entrepreneurs of marginally related ideas and products that “pile on” to existing echo chambers. From a risk-adjusted incentive standpoint, this would be a more rational approach for suppliers to take, than to try the higher-return but almost infinitely higher-risk alternative of creating a truly new idea. Overall, that doesn’t bode well for innovation, pluralism, and other liberal values.

Technical barriers, cultural barriers, differentiation

Naive globalization arguments tend to ignore the extent to which societies and individuals maintain independent desires and demands for boundaries. Some of this comes from traditional motivations behind trade protectionism, and some is cultural, emotional, and even religious.

When technology reduces or eliminates one barrier, other barriers rise in relative importance (sometimes even in absolute importance) as people intentionally find ways to “protect” and differentiate themselves from unrestricted flows. When container shipping made it economically viable to concentrate the production of skis in the most efficient American factories, the Japanese claimed that snow in Japan was distinctive and required a different length of ski that was only made in Japan. These kinds of nontariff barriers are common in culturally sensitive areas like food and the arts.²⁹ Nontariff barriers are often criticized (sometimes ridiculed) as indirect ways to subvert trade liberalization but that is a value judgment that assumes economic efficiency as the primary goal of human interaction, which of course it often isn’t. Minus any value assumptions, the question here is what kinds of new barriers will countries, firms, and people erect to counter some of the boundary-breaking consequences of machine translation in the next decade? There’s no way to fully anticipate the answer, but two kinds of barrier-raising seem almost inevitable because they directly connect to basic organizational and individual motivations.

Consider first the role of technical standards and particularly data structures as an indication of how firms might recreate barriers to competition (“moats”). Contemporary discussions about data portability notwithstanding, it is right now not straightforward to move a complex dataset from one CRM architecture to another; or from one cloud service to a competitors’ cloud service. When language barriers decline, the importance of these technical barriers rise in relative terms and might rise further through intentional action. Imagine then a world in which it is much easier to translate between Chinese and Russian than it is to “translate” between Salesforce and Oracle (really, between data stored in a Salesforce platform and an Oracle platform). That’s hardly a hypothetical given the challenges of interoperability at present; the interesting question is how much higher those boundaries rise. The same phenomenon should be expected for the internet generally, where interoperability is

²⁸ Axelrod (1997), 203–26.

²⁹ Aggarwal (2010), 221–44.

already declining for related reasons.³⁰ These already visible efforts are likely to be a taste of the creative experiments in barrier-raising that firms and governments develop over the next decade.

For individuals, the primordial desire to differentiate will probably yield even more creative strategies. One area where a significant demand for differentiation can be foreseen is in what the early-twenty-first-century “global elite” will use to mark itself off and qualify its own members. It’s interesting to note that the significance of language for elite differentiation is a modern phenomenon. The prebourgeois ruling classes of Europe were able to define themselves and cohere without a common language—when the King of England married a Spanish princess, they didn’t need to talk to each other very much. Noble marriage was a function of Machiavellian politics and shared kinship, and so an illiterate nobility could still function as a nobility in premodern times. But that’s not true for an industrial-era bourgeoisie: Their interactions and class consciousness depend in part on language, communication, and the economic and cultural coordination that follows. Durkheim’s concept of organic solidarity expresses the same mechanism through the division of labor.³¹ As Benedict Anderson wryly put it, you can sleep with anyone, but you can only read some peoples’ writing.³² And if you want to coordinate a complex division of labor (and rule or at least control it enough to be an effective capitalist) it helps enormously to have a common or at least a closely translatable language.

The ability to speak multiple languages fluently thus became a signal of elite differentiation across a broad swathe of the world. Anderson described it as a key point of connection between local colonial masters and the colonizing state, with bilingual elites in the colonies ruling monoglot populations and thus controlling the flow of information and authority from the metropole.³³ That function transformed during the postcolonial era into a meaningful signal of elite status (one among several of course). Poor and non-cosmopolitan commoners wouldn’t generally have the economic need to learn multiple languages, nor the resources to study them. And they wouldn’t get the value of cultural signaling that happens when you speak multiple languages in the first-class cabins of airplanes, which supports affiliative sorting in job and marriage markets.

When multiple languages were a luxury good that required significant investment of time and money to attain, the ability to speak and read them were signals that elites could use to identify each other. That signaling function has been in decline for a while (in part because of earlier technologies, which have made it easier and cheaper to learn languages). Machine translation will further reduce the signaling value toward a zero asymptote. Which—repeating the earlier point—will create a need and demand for other ways of signaling elite differentiation. There are plenty of possible paths for this demand to be expressed; precisely how that emerges is likely to be a bit of a surprise. It may very well turn out to be an even more exclusionary set of signals that are even harder for non-elites to attain, which (like some of the nontariff barriers mentioned previously) will tend to make the new borders *less* permeable than the old.

Amplify the upside

Machine translation could make the global economy more fractured and unequal over the next decade. But if we get a few important things right, machine translation could instead contribute to significant gains in economic productivity; to new and broader political and cultural coalitions; to further ethnic and genetic hybridization of the human species; to beneficial and sustainable forms of immigration; and to more profound cultural understandings that break down dysfunctional barriers limiting human progress.

³⁰Early results reported at Merrill and Weber (2020).

³¹Durkheim et. al (1933).

³²Anderson, B. (1983), 76.

³³Ibid., 83–111.

Productivity

Consider the possible productivity effects to start. The most straightforward mechanism is simply a next-generation globalization effect—as with container shipping in the twentieth century, boundaries (in this case, linguistic boundaries) between markets would fall leading to higher levels of competition and greater potential scale for production and distribution. The classical argument is that productivity gains emerge from the reallocation of resources across sectors following opening to trade. Other mechanisms that boost productivity include enhancing competition that puts pressure on domestic producers to lower price margins; improving efficiency and making greater efforts at innovation; and increasing the quality and variety of intermediate inputs that are available to domestic producers.³⁴

Studies of regional trade politics in the 1990s are one source of relevant models. Frankel et al. developed a gravity model of trade to test for the effect of regional trade agreements.³⁵ (Gravity models are built on the proposition that “baseline” trade between two countries is proportionate to the product of their GDPs and inversely proportionate to the physical distance between them.)³⁶ They found an independent effect of regional trade groupings that are layered “on top” of gravity model expectations, demonstrating that regional trade agreements (even though they often in the 1990s corresponded closely with physical contiguity) did more than simply place a label on trade flows that would have been present in a world without policy. But they also found an intriguingly large causal effect associated with common language. Incorporating a dummy variable for nine languages and comparing otherwise matched trade dyads that differ in language commonality, they found that two countries with strong linguistic ties tend to trade 65 percent more than two countries that have similar gravity model characteristics but different primary languages. This is a large effect and there are possible confounding variables—for example, countries with common languages frequently had colonial relationships in the past that may explain part of the effect through other mechanisms. But even if colonial links explain half the variance, there’s a significant effect for language commonality. And if machine translation were to quickly remove just one-half of the remaining language barrier effect between two markets, that would amount to roughly a 16 percent expected increase in trade—still a massive impact.

More recently, Joshi and Lahiri (2015) found that in cross-border research and development alliance formation some language friction may enable partners to rethink solutions and encourage collaboration but excessive friction can serve as an impediment on collaboration.³⁷ A 2018 study that makes use of a partial natural experiment broadly reinforces this estimate. Brynjolfsson et al. assess what happened when eBay introduced limited machine translation on its trading platform in 2014.³⁸ Looking at the impact of English–Spanish translation on US exports to Spanish-speaking Latin American countries the authors found an increase of 17–21 percent depending on the time windows in which the comparison is made. They found also that the increase is greater for differentiated products, products with more words in their listing titles, cheaper products, and less experienced buyers—which are consistent with the likely mechanism that machine translation is boosting trade by reducing search costs for buyers.³⁹

One shouldn’t make too much of the precise numbers that these models generate—they are rough estimates based on an imperfect baseline (gravity model) and an imperfect natural experiment (eBay). But the fact that the estimates fall in a similar range—and a range that is larger but not wildly inconsistent with some previous estimates of the degree to which language barriers can be trade inhibiting—is indicative of what machine translation will do to intensify trade, create larger more contestable

³⁴Helpman, E., P.R., Krugman. (1985); Aghion, P., N. Bloom, R. Blundell, R. Griffith and P. Howitt. (2005), 701–28; Grossman, G. and E. Helpman (1991), 43–61.

³⁵Frankel, J., E. Stein and S. Wei. (1995), 61–95.

³⁶By *baseline* I mean that gravity models suggest a “natural” economic determinant of trade, which is a useful baseline against which one can assess the impact of policy and other drivers.

³⁷Joshi, A. and N. Lahiri (2015), 123–52.

³⁸Brynjolfsson et al. (2018) NBER Working Paper 24917.

³⁹Put differently, products that have higher search costs that would be more significant in a buying decision tend to benefit more from machine translation. This finding increases confidence that the mechanism of trade facilitation is as expected.

markets, and reduce matching frictions *in well-resourced language dyads* at first.⁴⁰ The productivity boost could be significant, given recent IMF models estimating that a 1 percent absolute decline in tariffs can increase total factor productivity by as much as 2 percent.⁴¹

Even if that estimate is high, keep in mind that a small increase in productivity that continues over a meaningful period creates substantially greater absolute wealth levels. Productivity is colloquially the gift that keeps on giving, but only to those that experience its improvement in this case through the reduction of language barriers where machine translation works well.

Politics and culture

Coalition politics in domestic settings should be expected to shift as well, but with mixed consequences. At least two kinds of effects are foreseeable and though it's difficult to estimate magnitudes, they are still worth considering as logical mechanisms. The first would be a gradual but meaningful expansion of cross-linguistic and, by implication, cross-national political coalitions. From a "borderless world" perspective where barriers to movement of money and goods across national boundaries have fallen dramatically in the last 50 years, it's remarkable in relative terms how few truly cross-national political movements have emerged alongside economic globalization. Consider, for example, the outbreak of populist movements in the second half of the 2010s or anticapitalist movements (or at least antibank movements) in the first half of that decade. These were *multidomestic* phenomenon more than *transnational* ones, appearing in a number of countries simultaneously and for parallel reasons, but never really joining together to create a transnational movement.

Of course, language is not the only barrier here—political landscapes, concrete interests, cultural predilections, even electoral systems and rules also function as equivalents of nontariff barriers for politics. But language is likely a meaningful part of what holds back cross-national coalition formation and certainly impedes the mechanism of learning from parallel movements in other countries. A thought experiment that points in the right direction is to imagine an alternative history in which Marxist movements across Europe in the nineteenth through the twentieth centuries could communicate and interoperate without language friction. The parallel 2020s thought experiment might involve modern labor movements (particularly workers suffering the transnational shocks that will be associated with robotics), transnational climate coalitions, and perhaps transnational religious movements.

A second effect would be to modify what Robert Putnam called two-level games, where a leader appeals to a particular constituency at home and quite a different constituency abroad using distinct and sometimes incompatible arguments.⁴² A concrete example is former Israeli Prime Minister Benjamin Netanyahu, whose speeches in Hebrew for his domestic Israeli audience often had a very different tone and message than his speeches in English that were aimed at an international audience and often at the American Jewish community. Observers of Netanyahu's two-level game strategy often marvel at why he didn't suffer greater "hypocrisy costs" from these inconsistencies—and the language barrier is certainly part of the explanation.⁴³ As in some of the prior arguments, politicians wouldn't simply give up on the ability to play multiple games at once and would try to adapt by developing new ways to segregate messages—though none would probably be as simple and effective as language barriers.

⁴⁰See Lohmann, J. 2011. "Do Language Barriers Affect Trade?" *Economics Letters* 110: 159–62. Lohmann derives a Language Barrier Index, which is a measure of similarity between languages using categories from the World Atlas of Language. In his gravity model estimation, a 10 percent increase in the Language Barrier Index (which of course represents a 10 percent decrease in common linguistic features) is associated with a 7–10 percent decrease in trade flows between two countries.

⁴¹International Monetary Fund. 2016, March. "Reassessing the Productivity Gains from Trade Liberalization" (IMF Working Paper WP/16/77). Ahn, J., E. Dabla-Norris, R. A. Duval and B. Hu. A partially comparable study that assesses the (tricky) relationship between trade declines during the acute phase of the Global Financial Crisis and Great Recession around 2010 and declining productivity growth is Petersen Institute for International Economics. Hufbauer, G. C. and Z. Ju. 2016, October. *Increased Trade: A Key to Improving Productivity* (Issue Brief 16–15).

⁴²Putnam R.D. (1988), 427–60.

⁴³*Hypocrisy costs* is a loosely derived game-theoretic term to describe the reputation hit a leader supposedly would suffer from visible demonstrations of hypocritical behavior, making it harder for her to create credible commitments in the future.

A related and more macropolitical dynamic would likely emerge around exclusionary communities that are primarily defined by ethnicity. Language barriers can be instrumentally useful in keeping appeals and arguments largely contained within ethnically defined groups that speak the same language and that might not view those exclusionary arguments as racist, aggressive, or violence-inducing—while those outside the group probably would. As language boundaries fall and the range of constituents goes up, so does the range of opponents, enemies, and disruptors who have immediate access to the message. It's hard to foresee which directional effect predominates, but if you start with the assumption that “sunlight is a good disinfectant” for exclusionary appeals, machine translation is more likely to be a constraint on and net negative for harsh exclusionary arguments in ethnic politics.

The landscape for individuals choosing life partners and creating families would shift as well. It seems intuitive that linguistic homogamy (marriage among two people who speak the same language) is a desirable characteristic in a spouse, but this does limit substantially the size of de facto marriage markets even if you assume that these markets are mainly local or at best national (that has become a less robust assumption over the last decade as internet dating has exploded in popularity).⁴⁴ The percentage of interethnic marriages in the United States increased to 10.2 percent of households in 2016 from 7.4 percent in 2012, showing just how quickly behaviors can change as perceived barriers decline.⁴⁵ And while language obviously is not the only remaining barrier in marriage markets (ethnicity, nationality, religion, and culture intrude) a recent study from Switzerland (a multilingual country where there are both conationals and nonnationals, each with common or different languages) suggests that after spatial barriers, linguistic differences are the largest remaining obstacle to interethnic and international marriage.⁴⁶ A decline in linguistic homogamy would accelerate further the rise of interethnic marriages, whose precise effects on the further normalization of multiethnic children and other economic and sociocultural behaviors are almost certainly auspicious.⁴⁷

Another common life pattern that would be impacted is migration—particularly “voluntary” migrations that are less politically sensitive, such as during retirement. Somewhat like marriage markets, the geographic shape of demand for retirement migration is sensitive to language because retirees are less likely than young people to learn new languages even as they seek lower health care costs among other costs of living in other countries. Though the notion of “retiring abroad” has garnered popular attention in the United States over the last decade, it is still a rare decision, likely around 2 percent. And the most popular countries for retirement migration by Americans are those where either English or (in the cases of Mexico and Japan) the original birth language of the retiree is spoken.⁴⁸ As with other complex social dynamics, machine translation by itself won't impact the cultural, inertial, and other obstacles that restrain retirement migration but it would remove the language barrier—which could have meaningful short-term effects as well as contribute to longer-term reduction in more persistent barriers.

To try to make specific predictions about the magnitude of shifts in life patterns and cultural predilections would be reckless because language is only one component of these behaviors and is interdependent with other causal drivers. Still, this survey of some logical effects demonstrates the scope of what language barrier-breaking technologies could do to amplify the upside of liberal progress. It's important to keep in mind that these effects would concentrate among high-resource language

⁴⁴A 2019 paper estimates that 39 percent of heterosexual and 60 percent of same-sex American couples met online in 2017, up from 2 percent in 1995 and 22 percent in 2009. Rosenfeld, M., R. J. Thomas and S. Hausen (2019), 17753–58.

⁴⁵Data from US Census Bureau (<https://www.census.gov/library/stories/2018/07/interracial-marriages.html>) (accessed 27 August 27, 2021) shows the “barriers” that declined in this case are likely a combination of social norms and search costs (the same dynamic as internet dating noted previously), but not legal barriers (at least since the 1967 *Loving v. Virginia* case ruled antimiscegenation laws unconstitutional).

⁴⁶Followed by religious homogamy, age similarity, and educational status. See Schroedter, J. and J. Roessel. 2016, July. “The Importance of Linguistic Homogamy in (Inter)Marriages: Insights from a Multilingual Country.” Paper Presented at the 3rd ISA Forum of Sociology, Vienna, Austria.

⁴⁷See, e.g., the literature survey by Furtado, D. and S. Trejo (2012), 276–92. Edward Elgar Publishing.

⁴⁸See, e.g., table 5.J11 of the US Social Security Administration, Annual Statistical Supplement, 2018, <https://www.ssa.gov/policy/docs/statcomps/supplement/2018/5j.html#table5.j11>, measuring US citizens who receive social security abroad as about 1.5 percent. That is an imperfect proxy—e.g., a US retiree living abroad might maintain a bank account inside the United States and have her social security check deposited there.

dyads, and impact at a lower level in less well-resourced dyads. Poorly resourced dyads would see almost none of these effects in the short and medium term. They would then suffer in relative terms, as people and resources including goods and ideas are redirected toward interoperable language dyads (the equivalent of trade diversion effects on third parties that follow preferential tariff reductions).

That is a challenge that policy needs to address. As we argued in the introduction, there are distinct advantages to moving quickly and in anticipation of the deployment of next-generation machine translation systems. That argument should seem even more urgent after this survey of possible upside effects because many of them would tend toward positive feedback loops in which progress yields accelerating progress, at least for a while. The longer we wait to restructure incentives and practices around machine translation technology, the harder it will likely be to retrofit for broad liberal objectives.

What is to be done?

Consider from a design perspective how machine translation technology might be developed, licensed, distributed, and paid for. The leading platform firms at present have the largest incentives along with the greatest resources to make investments, but their incentives and business models will continue to skew toward high-resource language pairs. In addition, intelligence and security agencies of large governments will surely invest in a few less well-resourced languages of particular geopolitical interest to them. The irony is that medium and poorly resourced languages look set to be left behind *unless* the people and countries that speak them are perceived as a security risk by the United States and China in particular. This isn't a good outcome overall for human welfare in the next decade because it reinforces and could exacerbate existing inequalities both economic and cultural while undermining fairness and protection of the most vulnerable.

But technology needs to earn a return on investment. An obvious way to create greater balance among language dyads would be to subsidize the economic return on attention to low-resourced languages. Who might be motivated to fund such subsidies? Two possible models that suggest analogous rationales might be contemporary schemes that even out pharmaceutical investment and distribution and twentieth-century schemes that funded rural phone service. Machine translation does have some universal human rights-type characteristics, along the lines of access to small-molecule drugs that have large, fixed costs to discover but low variable costs to distribute and use. Even a minuscule surcharge on transactions made viable by machine translation in high-resource settings might be sufficient to fund such a scheme, without distorting the new value-creating possibilities that the technology will enable.

People who today argue that internet connectivity is a universal human right should consider whether language interoperability might be even more essential.⁴⁹ Language interoperability will be a prerequisite to exercising in meaningful ways the internet connectivity right that they are promoting. This is where the rationale behind subsidies for universal basic phone service in the twentieth century is a relevant analogy. On a global landscape even more so than a domestic one, the ability to join a network is worth very little if your language isn't functionally useful on the network. It might be a net negative if the only effective purpose of joining the network is to become a passive consumer of products (physical and cognitive) produced and distributed in other languages. It's a constructive stretch to imagine how existing international institutions like the Global Fund might extend their mission toward machine translation; or how the Global Fund model might be repurposed in a new institution combining governments and technology firms from the start.⁵⁰

Secondarily, it's important to consider *in advance* how to protect and compensate groups that come out as losers even in relative terms. The lessons of standard trade theory—and the failure of political elites to anticipate how those lessons would manifest in political economy during the last twenty years—are directly applicable. The populist backlash of the last decade didn't happen because political economy is silent on the challenge of relative losers from trade. The problem was that the presumed

⁴⁹Metzl (1996), 705–47.

⁵⁰See for a brief review: Feacham, R. and O. Sabot (2006), 537–40.

stabilization mechanism suggested by theory—using some of the surplus value created by opening trade to compensate losers and keep them on board politically—isn't natural or automatic and wasn't enacted by political elites. Without digressing into blame stories about the sources of that failure, there's an obvious lesson to be had about setting up of institutionalized mechanisms to compensate losers early in the process of technology development and rollout. It's reasonable to assume (though hard to measure precisely) and even more reasonable to act as if the costs of compensation rise, the longer the delay. In our view, political and ethical recriminations about elite failures in the 2000s and 2010s to act in the context of that prior era's globalization technologies, would be better redirected toward efforts to anticipate and act on machine translation consequences right now. A constructive mindset would focus on not making exactly the same mistakes again.

Another reason to do this in advance would be to reduce at least one of the motivations for security threats to machine translation technologies. Adversarial machine learning attacks on translation systems aren't the stuff of science fiction, and the consequences of even small manipulations of the underlying algorithms could be devastating and hard to detect until considerable damage is done.⁵¹ An adversarial machine learning attack on translation would be a new form of disinformation campaign, with a multiplier effect when combined with manipulated video and audio. Cybersecurity for machine translation systems is likely to become its own subspeciality in the security world, and the Luddite-type motivations of those who lose in relative terms is only one of many reasons why people might design attacks. But acknowledging that motivation and doing things to reduce it would also call attention to the needed upfront security investments aimed at hardening the underlying technologies and building awareness about threat and risk models in advance.

As in other cybersecurity and particularly adversarial machine learning issues, dealing with the technology will only go so far. Probably the most important and most challenging aspect for individuals and groups of people will be the risk of false positives in understanding, which we discussed in the "Challenges to global political economy" section of this article. Maintaining and extending diversity of thought and ideas in a world of widespread machine translation will be as important, though more abstract and harder to measure. The upsides of this technology revolution will be extraordinary and reveal themselves over time, but the transition period—which could be a decade or more—will have the most politically and socially salient and sensitive dilemmas to manage. Anticipation and preemptive management are by far the best strategy.

References

- Language Support [User Documentation]. <https://cloud.google.com/natural-language/docs/languages>
- Adler-Nissen, Rebecca and Alena Drieschova. 2019. "Track-Change Diplomacy: Technology, Affordances, and the Practice of International Negotiations." *International Studies Quarterly* 63(3): 531–45. <https://doi.org/10.1093/isq/sqz030>.
- Aggarwal, V. and S. Evenett. 2010. "Financial Crisis, 'New' Industrial Policy, and the Bite of Multilateral Trade Rules." *Asian Economic Policy Review* 5(2): 221–44.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith and P. Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *The Quarterly Journal of Economics* 120(2): 701–28
- Anderson, B. 1983. *Imagined Communities Reflections on the Origin and Spread of Nationalism*. Verso, 76.
- Anderson, Chris. 2006. *The Long Tail*. Hachette Books.
- Arivazhagan, N., A. Bapna, O. Firat, D. Lepikhin, M. Johnson, M. Krikun, ... Y. Wu. 2019. "Massively Multilingual Neural Machine Translation in the Wild: Findings and Challenges." *arXiv*. <https://arxiv.org/pdf/1907.05019.pdf> (accessed September 10, 2021).
- Awadalla, Hany. et al 2018. "Achieving Human Parity on Automatic Chinese to English News Translation." <https://www.microsoft.com/en-us/research/publication/achieving-human-parity-on-automatic-chinese-to-english-news-translation/> (accessed August 26, 2021).
- Axelrod, R. 1997. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41(2): 203–26.
- Barma, Naazneen, Durbin, Brent., and Weber, Steven. *How to Make the World More Liberal*. Forthcoming, Oxford University Press, 2022.
- Berk, Gerald. 1997. *Alternative Tracks: The Constitution of American Industrial Order 1865–1917*. Johns Hopkins University Press.
- Boulianne, Shelley and Yannis Theocharis. (2020). "Young People, Digital Media, and Engagement: A Meta-Analysis of Research." *Social Science Computer Review* 38(2): 111–27. <https://doi.org/10.1177/0894439318814190>

⁵¹A recent survey is Wang, X., J. Li, X. Kuang, Y. Tan, and Jin Li (2019), 12–23.

- Brooks, F. P. 1975. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley.
- Brynjolfsson, Erik, Yu (Jeffrey) Hu and Michael Smith. 2010. "The Longer Tail: The Changing Shape of Amazon's Sales Distribution Curve." <https://ssrn.com/abstract=1679991> (accessed September 18, 2021).
- Brynjolfsson, E., X. Hui and M. Liu. 2018. "Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform" (NBER Working Paper 24917).
- Brynjolfsson, Erik, Yu (Jeffrey) and Duncan Simester. 2011. "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." *Management Science* 57(8): 1373–86.
- Cettolo, M., C. Girardi, and M. Federico. 2012, May. "Wit3: Web Inventory of Transcribed and Translated Talks." In *EAMT 2012. Proceedings of the 16th EAMT Conference*, 261–68.
- Doyle, Michael W. 2016. "Liberalism and Foreign Policy." In *Foreign Policy: Theories, Actors, Cases*, edited by Steve Smith, Amelia Hadfield, and Timothy Dunne, ch. 3. Oxford University Press.
- Durkheim, E. and G. Simpson. 1933. *Émile Durkheim on the Division of Labor in Society*. Macmillan.
- Elberse, Anita and Felix Oberholzer-Gee. 2007. "Superstars and Underdogs: An Examination of the Long Tail Phenomenon in Video Sales." *MSI Reports: Working Paper Series* 4: 49–72 (Featured in HBS Working Knowledge).
- ELLORA. <https://www.microsoft.com/en-us/research/project/ellora/>
- Feacham, Richard and Oliver J. Sabot. 2006. "An Examination of the Global Fund at 5 Years." *The Lancet* 368(9534): 537–40.
- Frankel, J., E. Stein and S. Wei. 1995. "Trading Blocks and the Americas: The Natural, the Unnatural, and the Super-Natural." *Journal of Development Economics* 47: 61–95.
- Furtado, D. and S. Trejo. 2012. "Interethnic Marriages and Their Economic Effects." In *International Handbook on the Economics of Migration*, edited by A. F. Constant and K. F. Zimmerman, 276–92. Edward Elgar Publishing.
- Grossman, G. and E. Helpman. 1991. "Quality Ladders in the Theory of Growth." *The Review of Economic Studies* 58: 43–61.
- Helpman, E. and P. R. Krugman. 1985. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. MIT Press.
- Holtz, David, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer and Sinan Aral. 2020. "The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify." In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20)*. Association for Computing Machinery: 75–76.
- Interoperability. 1990. In *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*. IEEE.
- Joshi, Amol and Nandini Lahiri. 2015. "Language Friction and Partner Selection in Cross-Border R&D Alliance Formation." *Journal International Business Studies* 46: 123–52.
- Koehn, Phillip. 2005, September. "Europarl: A Parallel Corpus for Statistical Machine Translation." *MT Summit X. Proceedings of the 10th Machine Translation Summit*, 79–86. AAMT.
- Maimaiti, M Mieradilijiang, Yang Liu, Huanbo Luan, and Maosong Sun. 2019. "Multi-Round Transfer Learning for Low-Resource NMT Using Multiple High-Resource Languages." *ACM Transactions on Asian and Low-Resource Language Information Processing* 18(4): Article 38. Access date 28 November 2021. <https://doi.org/10.1145/3314945>.
- Islands, Mayr, E. 1959. "Isolation as an Evolutionary Factor." *Proceedings of the American Philosophical Society* 103(2): 221–30.
- Merrill, N. and S. Weber. 2020. "Web Site Blocking as a Proxy of Policy Alignment." *First Monday* 26(1). Access date 28 November 2021. <https://doi.org/10.5210/fm.v26i1.11415>.
- Meso, Peter and Radhika Jain. 2006. "Agile Software Development: Adaptive Systems Principles and Best Practices," *Information Systems Management* 23(3): 19. And access date can be 28 November 2021. <https://www.tandfonline.com/doi/abs/10.1201/1078.10580530/46108.23.3.20060601/93704.3>.
- Metzl, Jamie F. 1996. "Information Technology and Human Rights." *Human Rights Quarterly* 18: 705–47.
- Putnam, R. D. 1988. "Diplomacy and Domestic Politics: The Logic of Two-Level Games." *International Organization* 41(3): 427–60.
- Rosenfeld, M., R. J. Thomas and S. Hausen. 2019. "Disintermediating Your Friends: How Online Dating in the United States Displaces Other Ways of Meeting." *Proceedings of the National Academy of Sciences* 116(36): 17753–17758.
- Thompson, B. (2015) "Aggregation Theory." <https://stratechery.com/2015/aggregation-theory/> (accessed August 21, 2021).
- Tiedemann, Jörg. 2018. "Emerging Language Spaces Learned from Massively Multilingual Corpora." *DHN 2018. Proceedings of the 3rd Digital Humanities in the Nordic Countries*, 188–97. CEUR Workshop Proceedings.
- Wang, Xianmin, Jing Li, Xiaohui Kuang, Yu-an Tan and Jin Li. 2019. "The Security of Machine Learning in an Adversarial Setting: A Survey." *Journal of Parallel and Distributed Computing* 130: 12–23.
- Wang, Haifeng, Hua Wu, Zhongjun He, Liang Huang, and Kenneth Ward Church. 2021. *Progress in Machine Translation, Engineering*.
- Welch, H. Gilbert and William C. Black. 2010. "Overdiagnosis in Cancer," *JNCI: Journal of the National Cancer Institute* 102(9): 605–13. <https://doi.org/10.1093/jnci/djq099> (accessed August 26, 2021).
- Wills, A., Georgia Barrie, Jake Kendall. 2019. Conversational interfaces and the long tail of languages in developing countries The opportunities and challenges in the application of Natural Language Processing in low income countries Conversational interfaces.
- Ziemski, Michal, Marcin Junczys-Dowmunt and Bruno Pouliquen. 2016, May. "The United Nations Parallel Corpus v1.0." *LREC 2016. Proceedings of the 10th International Conference on Language Resources and Evaluation*, 3530–34. European Language Resources.