Language Barriers, Artificial Intelligence, and the Allocation of Talent

Joel Christoph*

Prepared for the NBER Economics of Talent Conference
Cambridge, MA, October 2, 2026

Abstract

Frontier knowledge in science, software, and entrepreneurship is documented predominantly in English. Able individuals whose strongest language is different face higher costs of accessing that knowledge, so latent talent in many language communities remains unused. This paper develops a theory of how multilingual artificial intelligence changes the identification, development, and allocation of talent by lowering language-dependent access costs. Individuals with heterogeneous ability choose between a local occupation and a frontier occupation whose returns depend on access to knowledge concentrated in a dominant language. Language distance creates community-specific ability thresholds for entry into frontier work, so equally able individuals sort differently across communities. Improvements in machine translation and multilingual model quality lower these thresholds. The core model yields four results. Participation in frontier fields rises with multilingual AI quality. The response is unambiguously larger in communities with greater language distance and in fields with greater dependence on the dominant language. Aggregate frontier output rises while average output per participant can fall, because entrants are drawn from the margin, so comparisons of averages understate the gains. As multilingual AI quality approaches its frontier, entry thresholds equalize across communities and the allocation converges to a language-blind benchmark in which ability alone determines selection. Three extensions study talent identification under noisy signals, the substitution between multilingual AI and skilled migration, and dynamic amplification through human capital investment. A final section shows that developers set language coverage to maximize product revenue rather than the social value of newly accessible talent, so equilibrium coverage is below the social optimum, which supports treating language coverage as talent infrastructure.

JEL codes: J24, O33, and F22.
Keywords: allocation of talent; language barriers; multilingual artificial intelligence.

1 Introduction

Research on talent has concentrated on ability, schooling, and local labor market conditions as the determinants of who reaches high achievement. This paper studies a further determinant: the language distance between a high-potential individual and the knowledge, mentors, and markets required to convert ability into output. Reference material, technical documentation, and elite professional networks in most frontier fields are concentrated in English. For an able person whose strongest language is different, this concentration acts as a

*FernUniversität in Hagen. Email: joel.christoph@proton.me. This project is supported by an Imminent Research Grant in Language Economics from Imminent, the research center of Translated. The author has no other sources of research support for this project and no financial relationships or other potential conflicts of interest to disclose under the NBER research disclosure policy.

tax on every step of talent development, from searching for technical answers to publishing work to entering global labor markets.

Recent progress in machine translation and multilingual large language models changes the size of this tax. Unlike schooling reforms or migration policy, the technology operates across many countries and career stages at low marginal cost. Its effect on the supply of talent is therefore a first-order question for the economics of talent, yet there is no formal framework that connects language technology to talent identification, development, and allocation. This paper provides such a framework.

The contribution is theoretical. I build a tractable occupational choice model in the Roy (1951) tradition in which language distance drives a wedge between latent ability and realized achievement. The wedge has a specific structure: it scales with the distance between an individual's strongest language and the dominant language of a field, with the degree to which the field's knowledge is concentrated in that dominant language, and inversely with the quality of multilingual AI. Comparative statics on AI quality then describe how the technology changes who participates in frontier work, who is identified as talented, who migrates, and how much of the world's latent ability is converted into output.

The core model produces four results. First, participation in frontier fields rises as multilingual AI improves, because the ability threshold for profitable entry falls. Second, the response is unambiguously larger in communities with greater language distance and in fields with greater dominant-language dependence, so the technology compresses cross-community participation gaps from below. Third, aggregate frontier output rises through an intensive margin and an extensive margin at once, while the average ability of participants falls, because entrants come from the margin. Evaluations that compare average productivity before and after an improvement will therefore understate the aggregate gain, and can even reverse its sign. Fourth, as AI quality approaches its frontier, entry thresholds equalize across language communities and the allocation converges to a language-blind benchmark in which selection depends on ability alone. Residual participation gaps at high AI quality then identify non-language barriers, so the model provides a decomposition of observed gaps into a language wedge and a residual wedge.

Three extensions widen the framework toward the central themes of the economics of talent. The first concerns identification. Gatekeepers such as admissions committees, referees, and employers observe ability through signals whose noise increases with language distance: an able applicant writing in a weaker language produces a less informative record. In a standard normal learning environment, this noise causes systematic under-identification of high-ability individuals from distant language communities. Multilingual AI raises signal precision, which raises the probability that a given able individual is identified. The mechanism formalizes one source of lost talent and predicts that AI reduces it. The second extension concerns migration. Before high-quality multilingual AI, the main private remedy for language distance was to migrate into the dominant-language environment. The model shows that AI and migration are substitutes: as AI quality rises, the ability cutoff above which migration is chosen also rises, so a growing share of frontier participants contribute from their home community. The composition of skilled migration then shifts toward the extreme top of the ability distribution, while sending communities retain more of their talent. The third extension adds a human capital investment stage. Because access costs also govern the cost of acquiring advanced skills, improvements in AI raise investment, so the long-run participation response exceeds the static response. Static evaluations understate the technology's effect on talent supply.

The final part of the paper asks whether markets will provide the right amount of language coverage.

Developers choose coverage to maximize product revenue, which reflects the current usage and purchasing power of each language community. The social value of coverage additionally includes the output of newly participating talent and the nonrival knowledge that talent produces, neither of which the developer appropriates. The communities where Proposition 2 locates the largest talent gains, those with high language distance and low current income, are the communities that private incentives serve last. Equilibrium coverage is therefore below the social optimum, with the largest gap where the social return is highest. The policy implication is that language coverage, translation quality for technical text, and multilingual evaluation benchmarks should be treated as talent infrastructure, comparable to scholarship systems or research capacity programs.

The paper proceeds as follows. Section 2 relates the framework to existing work. Section 3 presents the environment and Section 4 characterizes equilibrium participation. Section 5 contains the comparative statics on multilingual AI. Sections 6 through 8 develop the extensions on identification, migration, and dynamics. Section 9 studies the provision of language coverage. Section 10 states testable implications, Section 11 discusses policy, and Section 12 concludes. Proofs are collected in the Appendix.

2 Related Literature

The paper builds on four literatures. The first is the economics of the allocation of talent. Murphy et al. (1991) study how the sectoral allocation of the most able individuals affects growth, and Hsieh et al. (2019) show that declining barriers facing women and Black Americans account for a substantial share of United States growth since 1960. Bell et al. (2019) document that exposure to innovation strongly predicts who becomes an inventor, so equally able children face very different probabilities of reaching the innovation frontier. Agarwal and Gaule (2020) show that International Mathematical Olympiad medalists from poorer countries produce far less subsequent mathematics than equally ranked medalists from richer countries, direct evidence that latent talent at the very top goes unused. This paper contributes a barrier that has not been formalized in this literature: language distance to the dominant language of a field, together with a technology that compresses it. The misallocation perspective follows Restuccia and Rogerson (2017), applied to people rather than firms.

The second literature is the economics of language. Ginsburgh and Weber (2020) survey the field. Chiswick and Miller (2015) document large returns to dominant-language proficiency among migrants. In international trade, a common language raises bilateral trade substantially (Egger and Lassmann 2012; Melitz and Toubal 2014), and Brynjolfsson et al. (2019) show that the introduction of machine translation on a large digital platform raised exports along treated language pairs, causal evidence that AI-mediated language access changes economic behavior. This paper transfers that logic from goods to talent: the relevant flow is not merchandise, it is the participation of able individuals in frontier knowledge work.

The third literature concerns high-skilled migration. Borjas (1987) analyzes selection into migration, Kerr et al. (2016) survey global talent flows, and Docquier and Rapoport (2012) assess brain drain and development. The migration extension of this paper adds a technological substitute for migration and characterizes how it changes both the volume and the selection of talent flows.

The fourth literature is the economics of artificial intelligence (Agrawal et al. 2019). Recent evidence shows that generative AI raises productivity most for less experienced or lower-performing workers (Noy and Zhang 2023; Brynjolfsson et al. 2025). The present framework predicts an analogous pattern across language

communities: the largest gains accrue to those farthest from the dominant language. It also connects AI to the burden-of-knowledge problem identified by Jones (2009): if reaching the frontier requires absorbing ever more material, then technologies that lower the cost of absorbing material in any language expand who can reach it. Hjort and Poulsen (2019) show that the arrival of fast internet raised skilled employment in Africa; multilingual AI is a complementary access technology operating on the language margin rather than the connectivity margin.

3 Environment

There is a continuum of individuals partitioned into language communities indexed by ℓ. An individual is characterized by latent ability $a \geq 0$ with continuous distribution F and density f.

Assumption 1. *Latent ability is distributed identically across language communities.*

Assumption 1 isolates the allocation question: any difference in frontier participation across communities in this model reflects barriers rather than differences in talent. The assumption is deliberately strong. It corresponds to the working hypothesis behind the evidence in Agarwal and Gaule (2020), namely that raw ability at the top of the distribution appears in every country, while its conversion into output does not.

Each community has a language distance $d_\ell \geq 0$ to the dominant language of frontier knowledge, with $d_\ell = 0$ for the dominant-language community. Frontier occupations are grouped into fields indexed by f, and each field has a dependence parameter $E_f \in (0,1]$ measuring how concentrated its knowledge, tools, and networks are in the dominant language. Mathematics, with a compact symbolic core, has lower E_f than experimental biomedicine or entrepreneurship, where tacit knowledge and professional networks matter more. Multilingual AI quality is a scalar $T \in [0,1)$, where higher T means better translation, retrieval, and multilingual assistance. The effective access cost of a member of community ℓ entering field f is

$$c_{\ell f}(T) \;=\; E_f\, d_\ell\, (1-T).$$

Access costs rise with language distance and field dependence, and fall as multilingual AI improves. When $T \to 1$, access costs vanish for every community.

An individual chooses between a local occupation paying wa, with $w > 0$, and a frontier occupation in field f. Frontier output is complementary between ability and access to knowledge: gross frontier earnings are $\pi\, a\, e^{-c_{\ell f}(T)}$ with $\pi > w$, net of a fixed entry cost $k > 0$ covering learning, credentialing, and integration into the field. An individual therefore enters the frontier occupation if and only if

$$\pi\, a\, e^{-c_{\ell f}(T)} - k \;\geq\; wa.$$

The multiplicative form embodies the central complementarity: access to frontier knowledge is more valuable for more able individuals, so barriers to access distort the allocation of the most able most severely. Earnings here should be read broadly as the private return to frontier work, including career and status returns, while the fixed cost should be read as the lump of investment required before any frontier output can be produced.

4 Equilibrium Participation and the Language Wedge

Define the net return factor $\Delta(c) = \pi e^{-c} - w$. If $\Delta(c_{\ell f}) \leq 0$, no member of community ℓ enters field f: the language barrier is so large that the frontier occupation is dominated at every ability level. If $\Delta(c_{\ell f}) > 0$, entry follows a cutoff rule. The marginal entrant has ability

$$a^*_{\ell f}(T) \;=\; \frac{k}{\Delta\big(c_{\ell f}(T)\big)} \;=\; \frac{k}{\pi e^{-E_f d_\ell (1-T)} - w},$$

and the participation rate of community ℓ in field f is $1 - F\big(a^*_{\ell f}(T)\big)$.

Two properties of the cutoff organize the analysis. First, $a^*_{\ell f}$ is increasing in the composite barrier $b_{\ell f} \equiv E_f d_\ell$: communities farther from the dominant language, and fields more dependent on it, require higher ability for entry. Second, in the dominant-language community the cutoff is $\bar{a}^* = k/(\pi - w)$ regardless of T. The interval $[\bar{a}^*, a^*_{\ell f}(T))$ therefore contains individuals who would enter frontier work if born into the dominant-language community but do not enter from community ℓ. Their mass,

$$M_{\ell f}(T) \;=\; F\big(a^*_{\ell f}(T)\big) - F(\bar{a}^*),$$

measures misallocated talent attributable to language alone. The object $M_{\ell f}$ is the model's counterpart of the lost inventors in Bell et al. (2019) and the unused Olympiad talent in Agarwal and Gaule (2020), with a language interpretation.

5 Multilingual AI and the Allocation of Talent

Proposition 1 (Participation)**.** *Wherever entry is viable, the cutoff $a^*_{\ell f}(T)$ is strictly decreasing in T, and the participation rate $1 - F(a^*_{\ell f}(T))$ is strictly increasing in T.*

Proposition 2 (Heterogeneous gains)**.** *Conditional on viability, the magnitude of the participation response, $\left|\partial a^*_{\ell f}/\partial T\right|$, is strictly increasing in the composite barrier $b_{\ell f} = E_f d_\ell$. The largest gains from multilingual AI accrue to communities with greater language distance and to fields with greater dominant-language dependence.*

Proposition 3 (Extensive margin at the viability boundary)**.** *For a community-field pair with $\Delta(c_{\ell f}(T)) \leq 0$, participation is zero. There exists a threshold quality $T^\dagger_{\ell f} = 1 - \ln(\pi/w)/(E_f d_\ell)$ such that participation becomes positive for all $T > T^\dagger_{\ell f}$. The threshold is increasing in the composite barrier, so the most excluded communities are activated last, and each activation is a discrete event in which participation jumps from zero to a positive mass.*

Propositions 1 through 3 jointly describe the incidence of the technology. Communities close to the dominant language gain little because they had little to gain. Communities at intermediate distance gain along a smooth intensive margin of falling cutoffs. Communities beyond the viability boundary gain nothing until quality crosses their activation threshold, at which point they gain discretely. Cross-community participation gaps therefore close from below, and the time path of closing is convex in language distance.

Proposition 4 (Composition)**.** *Aggregate frontier output of community ℓ in field f,*

$$Q_{\ell f}(T) = \int_{a^*_{\ell f}(T)}^{\infty} \pi a e^{-c_{\ell f}(T)} dF(a),$$

is strictly increasing in T, through an intensive-margin channel (incumbent participants become more productive as access costs fall) and an extensive-margin channel (new entrants join). Average ability among participants, $\mathbb{E}\Big[a \mid a \geq a^*_{\ell f}(T)\Big]$, is strictly decreasing in T.

Proposition 4 has a measurement implication. Because entrants are drawn from the margin, average output per participant can fall while total output rises. Evaluations of multilingual AI that compare the average productivity of participants before and after an improvement will understate, and can even reverse the sign of, the aggregate talent gain. The correct object is the total output of the community, including the output of those who would otherwise not have participated at all. The same logic applies to selective institutions: a program admitting from a newly activated language community should expect its marginal admits to look weaker on average even when the program's total contribution to output has risen.

Proposition 5 (Convergence to the language-blind benchmark)**.** *As $T \to 1$, $c_{\ell f}(T) \to 0$ for every community and field, cutoffs converge to $\bar{a}^* = k/(\pi - w)$, misallocated mass $M_{\ell f}(T) \to 0$, and the allocation of talent converges to a benchmark in which selection into frontier work depends on ability alone.*

The benchmark clarifies what multilingual AI can and cannot do. It equalizes the language component of access costs, so residual cross-community differences in participation at high T identify non-language barriers such as capital constraints, discrimination, or weak local institutions. The model thus provides a decomposition: observed participation gaps combine a language wedge, which the technology closes, with other wedges, which it does not. This decomposition matters for interpretation. If participation gaps persist after language costs have been compressed, the remaining gaps should not be attributed to language, and policy attention should turn to the residual barriers.

6 Talent Identification under Noisy Signals

The core model treats ability as privately known and entry as an individual choice. Much of the identification of talent, however, is performed by gatekeepers: admissions committees, scholarship panels, journal referees, and employers. This section shows that language distance degrades the signals on which gatekeepers rely, and that multilingual AI improves identification by raising signal precision.

Ability is drawn from a normal prior, $a \sim N(\mu, \tau^2)$. A gatekeeper observes the signal

$$s \;=\; a + \sigma_{\ell f}(T)\,\varepsilon, \qquad \varepsilon \sim N(0,1), \qquad \sigma_{\ell f}(T) \;=\; \sigma_0 + \eta\, E_f\, d_\ell\, (1-T),$$

with $\sigma_0 > 0$ and $\eta > 0$. The specification states that an application, manuscript, or interview produced across a language barrier is a noisier measure of underlying ability, and that multilingual AI reduces the added noise. The gatekeeper admits the candidate when the posterior mean of ability exceeds a standard $\bar{q} > \mu$, which is equivalent to admitting when s exceeds a cutoff $s^*(\sigma) = \mu + (\bar{q} - \mu)(\tau^2 + \sigma^2)/\tau^2$, increasing in σ.

Proposition 6 (Identification)**.** *Fix an ability level $a > \mu + (\bar{q} - \mu)(\tau^2 - \sigma^2)/\tau^2$. The probability that an individual of ability a is admitted, $\Phi\big((a - s^*(\sigma))/\sigma\big)$, is strictly decreasing in σ and therefore strictly increasing in T. High-ability individuals from distant language communities are systematically under-identified, and improvements in multilingual AI reduce the under-identification.*

The mechanism operates through Bayesian shrinkage. Noisier signals are shrunk more heavily toward the prior mean, so an outstanding candidate whose record is filtered through a language barrier looks merely good, while the admission standard applied to the noisy signal rises. Both forces work against able candidates from distant communities. The proposition implies that multilingual AI functions as an identification technology, not only as a production technology: it changes who is seen, in addition to changing what those who enter can produce. It also implies a bias in retrospective evaluations of talent-search programs. Programs that screened across language barriers before high-quality multilingual AI will have rejected a disproportionate share of their strongest distant-language applicants, so their historical admit data understate the ability available in those communities.

7 Migration and Remote Participation

Before high-quality multilingual AI, the main private remedy for language distance was migration into the dominant-language environment. This section adds that option and characterizes the substitution between the two technologies of access. An individual in community ℓ now has three options: the local occupation paying wa; frontier work from home paying $\pi a e^{-c_{\ell f}(T)} - k$; or migration at cost $m > 0$, which removes the language penalty and pays $\pi a - k - m$. The gain from migrating relative to frontier work from home, $\pi a(1 - e^{-c}) - m$, is increasing in ability, so migration is chosen by the ablest.

Proposition 7 (Substitution between AI and migration)**.** *Among frontier participants, migration is chosen by individuals with ability above*

$$a^M_{\ell f}(T) \;=\; \frac{m}{\pi\big(1 - e^{-c_{\ell f}(T)}\big)},$$

which is strictly increasing in T. As multilingual AI improves, the migration cutoff rises, the share of frontier participants who migrate falls, and migrants are drawn from an increasingly extreme upper tail of the ability distribution.

Together with Proposition 1, this yields a distinctive joint prediction: total frontier participation from a distant language community rises with AI quality while the migrant share of that participation falls. Sending communities retain more of their talent, in the sense of Docquier and Rapoport (2012), without losing frontier output, because retained individuals now contribute remotely. The selection result also revises the classic analysis of Borjas (1987): the degree of positive selection into skilled migration is not a fixed feature of sending countries, it is decreasing in the quality of the technologies that substitute for physical relocation. For destination-country institutions that recruit talent, the model predicts a compositional shift in applicant pools toward the extreme top, alongside growth of remote collaboration with non-migrants.

8 Dynamic Amplification through Human Capital Investment

Access costs govern more than the return to participating today. They also govern the cost of building the skills that determine tomorrow's ability to participate, because textbooks, courses, and mentorship in frontier fields are concentrated in the dominant language. This section adds an investment stage and shows that it amplifies the static effects.

Life has two periods. In period 1, an individual may invest in advanced human capital at cost $\chi\, e^{c_{\ell f}(T)}$, with $\chi > 0$: the same language barrier that taxes participation taxes preparation, at an exponential rate consistent with the production side. Investment raises period-2 frontier productivity from $\pi a e^{-c}$ to $(1+\gamma)\pi a e^{-c}$, with $\gamma > 0$, and does not affect the local wage. Individuals decide on investment in period 1 and on occupation in period 2.

Corollary 1 (Amplification)**.** *Among period-2 frontier participants, investment is chosen by individuals with ability above $a^I_{\ell f}(T) = \chi e^{c}/(\gamma\pi e^{-c}) = \chi e^{2c_{\ell f}(T)}/(\gamma\pi)$, which is strictly decreasing in T at twice the exponential rate of the production cost. An improvement in multilingual AI therefore raises frontier output through three channels at once: incumbent productivity, new entry, and increased skill investment among both groups. The long-run response of output to T strictly exceeds the static response of Proposition 4.*

The doubling of the exponent is not an artifact: the investing individual crosses the language barrier twice, once to acquire the skill and once to apply it. The corollary implies that short-run evaluations of multilingual AI understate its effect on talent supply, since the investment channel takes a schooling cycle or a training cycle to appear in output. It also identifies where educational complements matter most: in communities near the investment margin, pairing language technology with instruction converts the technology's cost reduction into human capital rather than only into current output.

9 Endogenous Language Coverage and Underprovision

The analysis so far treats T as common across languages. In practice, developers choose how much quality to provide per language, and quality differs widely across languages. This section asks whether private provision delivers the socially efficient pattern of coverage.

Let a developer choose community-specific quality $T_\ell \in [0,1)$ at convex cost $\kappa(T_\ell)$, with $\kappa' > 0$ and $\kappa'' > 0$. Private revenue from community ℓ is $R_\ell(T_\ell) = \rho\, y_\ell\, u_\ell\, T_\ell$: proportional to the community's current income per user y_ℓ, its current user base u_ℓ, and delivered quality. The social value of quality adds the talent output created by coverage, $V_\ell(T_\ell) = R_\ell(T_\ell) + \sum_f \lambda_f\, Q_{\ell f}(T_\ell)$, where $\lambda_f > 0$ converts frontier output into social value and includes the nonrival component of knowledge production that no participant, and no developer, appropriates.

Proposition 8 (Underprovision)**.** *The developer's optimum satisfies $\kappa'(T^m_\ell) = \rho\, y_\ell\, u_\ell$, while the planner's optimum satisfies $\kappa'(T^p_\ell) = \rho\, y_\ell\, u_\ell + \sum_f \lambda_f\, \partial Q_{\ell f}/\partial T$. Since $\partial Q_{\ell f}/\partial T > 0$ wherever participation is positive or attainable, $T^m_\ell < T^p_\ell$ for every such community. The gap $T^p_\ell - T^m_\ell$ is largest where the talent term is large relative to the revenue term: communities with high language distance, low current income, and substantial latent-ability mass near the entry margin.*

The proposition inverts a common intuition about market prioritization. Revenue-maximizing coverage follows current purchasing power, which is lowest in the communities where, by Propositions 2 and 3, the marginal social product of coverage is highest. Language coverage for technical and scientific text is therefore an underprovided input into the global production of talent, in the same sense in which basic research is underprovided. The natural remedies are the standard ones for public-good underprovision: public or philanthropic funding of open multilingual models and benchmarks for underserved languages, procurement of technical-domain translation quality, and support for the evaluation infrastructure that makes quality in a given language measurable at all. The framework supplies the targeting rule: fund the languages with the largest product of distance, latent-talent mass at the margin, and field dependence, which is where private incentives are weakest and social returns are strongest.

10 Testable Implications

The theory yields predictions that distinguish it from alternatives and structure future empirical work. First, improvements in multilingual AI quality for a given language should raise frontier participation from that language community, with larger responses where baseline distance to the dominant language is greater (Propositions 1 and 2), and with discrete activations of previously absent communities (Proposition 3). Second, within a community, responses should be larger in fields with greater dominant-language dependence than in fields whose knowledge is already available locally, a difference-in-differences structure across fields within languages. Third, entrant cohorts following large quality improvements should have lower average observable ability proxies than incumbent cohorts while total community output rises (Proposition 4). Fourth, acceptance rates of high-ability candidates from distant language communities at selective institutions should rise with multilingual AI quality, conditional on underlying ability proxies determined before the improvement (Proposition 6). Fifth, the migrant share of a community's frontier participants should fall as its language coverage improves, while the average selectivity of those who still migrate rises (Proposition 7). The author's ongoing measurement project on the economics of multilingual AI, which constructs language-level indices of AI quality and inclusion, is designed to take these predictions to data in subsequent work.

11 Policy Implications

Three implications follow for institutions concerned with talent. For talent-search and scholarship programs, Proposition 6 implies that historical screening across language barriers rejected a disproportionate share of the strongest distant-language applicants; re-weighting outreach toward recently activated language communities is the corrective with the highest expected yield. For science and innovation policy, Proposition 8 implies that multilingual coverage of technical text is talent infrastructure whose private provision is weakest where its social return is highest, so public and philanthropic funders should treat language coverage, domain-specific translation quality, and multilingual evaluation benchmarks as fundable research inputs. For the interpretation of cross-country statistics, Propositions 4 and 5 counsel against reading average-productivity comparisons or persistent participation gaps at face value: the first understates gains from access technologies, while the second conflates language wedges with residual barriers that require different instruments.

12 Conclusion

Talent is misallocated when equally able people face unequal access to the knowledge that converts ability into achievement. Language is one of the oldest such barriers and, with current AI systems, one of the most compressible. The framework developed here shows that the resulting gains are largest precisely where they are least visible in average productivity statistics; that the technology improves the identification of talent, not only its productivity; that it substitutes for migration while raising the selectivity of those who still migrate; that its long-run effects exceed its short-run effects through the human capital channel; and that private provision of language coverage will not deliver these gains on its own. For a research program concerned with barriers that prevent high-potential individuals from realizing their potential, language distance deserves a place alongside credit constraints, information frictions, and discrimination.

Acknowledgements and Disclosures

This project is supported by an Imminent Research Grant in Language Economics awarded by Imminent, the research center of Translated. The author declares no other sources of financial support for this research and no financial relationships or other potential conflicts of interest relevant to this submission, in accordance with the NBER research disclosure policy.

A Proofs

Proof of Proposition 1. $\partial c_{\ell f}/\partial T = -E_f d_\ell < 0$ for $d_\ell > 0$, and $a^* = k/\Delta(c)$ with $\Delta'(c) = -\pi e^{-c} < 0$, so $\partial a^*/\partial T = -k\pi e^{-c} E_f d_\ell/\Delta(c)^2 < 0$. Participation $1 - F(a^*)$ is then strictly increasing in T since F is strictly increasing. □

Proof of Proposition 2. Write $|\partial a^*/\partial T| = k\pi b\, e^{-c}/\Delta(c)^2$ with $c = b(1-T)$ and $b = E_f d_\ell$. Differentiating in b gives a first term $k\pi e^{-c}/\Delta^2 > 0$ and a second term proportional to

$$\frac{d}{dc}\left[\frac{e^{-c}}{\Delta(c)^2}\right] = \frac{-e^{-c}\Delta^2 + 2e^{-c}\Delta\,\pi e^{-c}}{\Delta^4} = \frac{e^{-c}\big(2\pi e^{-c} - \Delta\big)}{\Delta^3} = \frac{e^{-c}\big(\pi e^{-c} + w\big)}{\Delta^3},$$

which is strictly positive whenever $\Delta > 0$. Both terms are positive, so the response is strictly increasing in b. □

Proof of Proposition 3. $\Delta(c) > 0$ if and only if $c < \ln(\pi/w)$, that is, $E_f d_\ell(1-T) < \ln(\pi/w)$, equivalent to $T > T^\dagger_{\ell f} = 1 - \ln(\pi/w)/(E_f d_\ell)$ when $E_f d_\ell > \ln(\pi/w)$. The threshold is increasing in $E_f d_\ell$. At $T = T^\dagger_{\ell f}$, $a^* \to \infty$ and participation is zero; for $T > T^\dagger_{\ell f}$, a^* is finite, so participation is positive whenever F places mass above a^*. □

Proof of Proposition 4. Differentiating $Q_{\ell f}(T)$,

$$\frac{\partial Q_{\ell f}}{\partial T} = \underbrace{E_f d_\ell \int_{a^*}^{\infty} \pi a e^{-c}\, dF(a)}_{\text{intensive margin} > 0} + \underbrace{\big(-\partial a^*/\partial T\big)\,\pi a^* e^{-c} f(a^*)}_{\text{extensive margin} > 0} > 0.$$

Average ability among participants is $\mathbb{E}[a \mid a \geq a^*]$, strictly increasing in a^* for any continuous F with positive density at a^*; since a^* is strictly decreasing in T, the average is strictly decreasing in T. □

Proof of Proposition 5. Immediate from $c_{\ell f}(T) = E_f d_\ell (1-T) \to 0$ as $T \to 1$, so $a^*_{\ell f} \to k/(\pi - w) = \bar{a}^*$ and $M_{\ell f} \to 0$. □

Proof of Proposition 6. With normal prior and normal noise, the posterior mean is $\mu + \frac{\tau^2}{\tau^2+\sigma^2}(s-\mu)$, so admission occurs when $s \geq s^*(\sigma) = \mu + (\bar{q}-\mu)\frac{\tau^2+\sigma^2}{\tau^2}$. Conditional on ability a, $s \sim N(a, \sigma^2)$, so the admission probability is $P(a,\sigma) = \Phi\big((a - s^*(\sigma))/\sigma\big)$. Differentiating the argument in σ,

$$\frac{d}{d\sigma}\left[\frac{a-s^*(\sigma)}{\sigma}\right] = \frac{-\frac{2\sigma(\bar{q}-\mu)}{\tau^2}\,\sigma - \big(a - s^*(\sigma)\big)}{\sigma^2} = -\frac{1}{\sigma^2}\left[a - \mu - (\bar{q}-\mu)\frac{\tau^2-\sigma^2}{\tau^2}\right],$$

which is strictly negative for all $a > \mu + (\bar{q}-\mu)(\tau^2-\sigma^2)/\tau^2$. For such a, $P(a,\sigma)$ is strictly decreasing in σ; since $\sigma_{\ell f}(T)$ is strictly decreasing in T, P is strictly increasing in T. □

Proof of Proposition 7. The migration gain relative to home frontier work is $G(a) = \pi a(1-e^{-c}) - m$, strictly increasing in a, zero at $a^M = m/[\pi(1-e^{-c})]$. Since c is strictly decreasing in T, $1-e^{-c}$ is strictly decreasing in T, so a^M is strictly increasing in T. The migrant share among participants, $[1-F(\max\{a^M, a^*\})]/[1-F(a^*)]$, then falls as a^M rises and a^* falls. □

Proof of Corollary 1 and Proposition 8. A period-2 frontier participant invests when $\gamma\pi a e^{-c} \geq \chi e^{c}$, giving $a^I = \chi e^{2c}/(\gamma\pi)$, strictly decreasing in T. Output then gains an investment term in addition to the two margins of Proposition 4, so the total derivative in T strictly exceeds the static derivative. For Proposition 8, the first-order conditions are as stated in the text; $\partial Q_{\ell f}/\partial T > 0$ from Proposition 4 (or the activation of Proposition 3), so $\kappa'(T^p_\ell) > \kappa'(T^m_\ell)$, and convexity of κ gives $T^p_\ell > T^m_\ell$. The comparative statics of the gap follow from the monotonicity of $\partial Q_{\ell f}/\partial T$ in the composite barrier (Proposition 2) and from the revenue term's proportionality to current income and usage. □

References

Agarwal, Ruchir, and Patrick Gaule. 2020. "Invisible Geniuses: Could the Knowledge Frontier Advance Faster?" *American Economic Review: Insights* 2(4): 409–424.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, eds. 2019. *The Economics of Artificial Intelligence: An Agenda*. Chicago: University of Chicago Press.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *Quarterly Journal of Economics* 134(2): 647–713.

Borjas, George J. 1987. "Self-Selection and the Earnings of Immigrants." *American Economic Review* 77(4): 531–553.

Brynjolfsson, Erik, Xiang Hui, and Meng Liu. 2019. "Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform." *Management Science* 65(12): 5449–5460.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *Quarterly Journal of Economics* 140(2): 889–942.

Chiswick, Barry R., and Paul W. Miller. 2015. "International Migration and the Economics of Language." In *Handbook of the Economics of International Migration*, Vol. 1, edited by Barry R. Chiswick and Paul W. Miller. Amsterdam: Elsevier.

Docquier, Frédéric, and Hillel Rapoport. 2012. "Globalization, Brain Drain, and Development." *Journal of Economic Literature* 50(3): 681–730.

Egger, Peter H., and Andrea Lassmann. 2012. "The Language Effect in International Trade: A Meta-Analysis." *Economics Letters* 116(2): 221–224.

Ginsburgh, Victor, and Shlomo Weber. 2020. "The Economics of Language." *Journal of Economic Literature* 58(2): 348–404.

Hjort, Jonas, and Jonas Poulsen. 2019. "The Arrival of Fast Internet and Employment in Africa." *American Economic Review* 109(3): 1032–1079.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87(5): 1439–1474.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies* 76(1): 283–317.

Kerr, Sari Pekkala, William Kerr, Çağlar Özden, and Christopher Parsons. 2016. "Global Talent Flows." *Journal of Economic Perspectives* 30(4): 83–106.

Melitz, Jacques, and Farid Toubal. 2014. "Native Language, Spoken Language, Translation and Trade." *Journal of International Economics* 93(2): 351–363.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1991. "The Allocation of Talent: Implications for Growth." *Quarterly Journal of Economics* 106(2): 503–530.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381(6654): 187–192.

Restuccia, Diego, and Richard Rogerson. 2017. "The Causes and Costs of Misallocation." *Journal of Economic Perspectives* 31(3): 151–174.

Roy, Andrew D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3(2): 135–146.