

Intelligence Hours

Local Prices and the Labor Time Cost of Paid AI Access

Kwangseob Ahn

School of Business, Sejong University, Republic of Korea

INLEVEL9, Republic of Korea

Revised public paper · Version 1.1 · 29 September 2026

Correspondence: about@Intelligence-hours.com

The initial manuscript completed peer review in private circulation before public release, as confirmed by the author. This revision responds to subsequent critical feedback. That review history is distinct from journal acceptance and from independent reproduction of the empirical material.

Abstract

Comparing the time needed to buy paid artificial intelligence services requires specifying the local price, product, worker population and wage jurisdiction. This article develops Intelligence Hours (IH) as a measurement framework whose primary estimand is the verified local monthly checkout price divided by the applicable gross hourly wage. It distinguishes that measure from a posted-price approximation and a common US-dollar benchmark. With a common positive price, the latter's ranking is exactly the inverse of dollar-converted wages; it therefore contributes an interpretable unit, not independent evidence of international AI-price differences. The framework separates matched-plan affordability from the lowest verified paid-entry cost within a defined catalog, makes geographic coverage part of each reported result, and specifies missingness, vintage and rank-sensitivity rules. A documentary illustration uses Claude Pro's September 2023 US and UK launch prices with official wage floors. The posted-price ratio is 1.176 hours for Washington, DC, 2.759 at the US federal benchmark and 1.727 at the UK adult rate. Tax coverage is unresolved, so none is represented as a verified checkout observation or a national welfare ranking. A separately reproduced Korean series illustrates the common-price accounting baseline. The original cross-country panel is not used to establish new empirical findings. The article provides an implementable protocol for local-price research while identifying the transaction evidence and coverage needed before publishing a global local-price index.

Keywords: artificial intelligence; local prices; affordability; minimum wages; wage jurisdictions; labor time

JEL classification: C81; J38; L86; O33

1 Introduction

A monthly AI subscription can require very different amounts of work depending on the price available to a purchaser and the wage applicable to that worker. Expressing the cost in hours is immediately interpretable: 2.5 hours means two hours and thirty minutes, without an arbitrary index base. Interpretation alone, however, does not establish that a country ranking measures differences in AI markets. If every country is assigned the same dollar price, the ranking contains exactly the same ordering information as dollar-converted hourly wages.

This observation changes the research question. The relevant question for paid access is how observed local prices, product eligibility and applicable wages jointly determine the time cost faced by a specified population. A common-price calculation remains useful as a transparent purchasing-power benchmark, but it cannot reveal geographic price differentiation that it has excluded by construction. The primary measure in this article is therefore local-price IH; the common-dollar series is a comparison baseline.

Three questions organize the analysis. First, what constitutes an observed local subscription price when taxes, billing channels and eligibility differ? Second, how should an indicator represent wages when statutory floors vary within countries? Third, how can one compare access to a named service separately from the least expensive qualifying paid option? These questions require different price catalogs and explicit reporting units. They cannot be resolved by adding more products to a common-dollar median.

The paper contributes an operational measurement protocol and a bounded documentary demonstration. It does not claim to invent price-to-wage ratios, identify a new latent measure of intelligence, or establish a causal explanation for adoption. The empirical demonstration is deliberately narrow: historical posted prices and wage floors can be checked, while the tax-inclusive transaction prices needed for the principal local measure remain unavailable. Defining an estimand and implementing a complete global dataset are separate accomplishments.

2 Related research and the measurement problem

Labor-time prices have a substantial precedent. Ashenfelter (2012) studies real wages using workers performing standardized jobs, while the Big Mac literature studies what a recognizable product can reveal about international prices and non-traded inputs (Parsley & Wei, 2007). IH retains the readability of time prices, but statutory wage floors do not standardize jobs, productivity or the workers represented. A familiar subscription also does not remain a fixed-quality product when models, quotas and bundled services change.

Affordability research distinguishes a posted charge from the resources left after essential expenditure. Barrantes and Galperin (2008), Hilbert (2010) and Stone (2006) motivate attention to income distributions, budgets and definitions of need. A ratio using gross statutory pay does not observe taxes, transfers, living costs, wage enforcement or available hours of employment. Legal coverage and effective pay also differ (Kanbur & Ronconi, 2018). Accordingly, IH describes a specified price burden, not the proportion of people able to afford the service.

The wider AI-access literature reinforces this boundary. Humlum and Vestergaard (2025) document unequal adoption among workers; Ahia et al. (2023) examine language-related differences in API token costs and utility. Neither supplies a conversion from subscription affordability to effective access. Productivity evidence from a particular workplace likewise cannot be applied as a universal multiplier for labor saved (Brynjolfsson et al., 2025). Task-based automation models require output, quality and production costs that a monthly access fee does not measure (Zeira, 1998; Acemoglu & Restrepo, 2019).

The design below follows the price-measurement emphasis of Cavallo and Rigobon (2016): the product, market, observation date and collection method are part of the observation. Lower-priced tiers must be recorded with their actual restrictions. Quality-differentiation theory offers possible explanations for tiers, but a price list alone does not identify a provider's motive or establish equivalence between plans (Deneckere & McAfee, 1996).

3 Three price concepts and two access questions

3.1 The primary local checkout measure

Let j identify a wage jurisdiction and stated worker population, p a plan version and t a reference date. Let P denote the verified recurring monthly checkout charge in local currency for an eligible purchaser in that jurisdiction, and W the applicable gross hourly wage in the same currency. The primary measure is

$$IH_{jpt}^{local} = \frac{P_{jpt}}{W_{jt}} \quad (1)$$

The charge includes mandatory taxes and unavoidable provider checkout fees where they apply. A card issuer's foreign-exchange spread or account-specific payment fee belongs in a separately specified payment scenario unless directly observed for the purchaser represented. A cross-border dollar charge is not converted into a verified local checkout price merely by multiplying it by a central-bank exchange rate.

A value can be published under Equation 1 only when market, currency, billing period, tax treatment, eligibility and effective date are established. If the price is unavailable or tax treatment is unresolved, the primary result is unavailable. The measure remains gross-wage equivalent hours, not hours of disposable income or actual labor supplied. Free services, devices, connectivity, payment access and the ability to use the product require separate indicators.

3.2 Posted prices and a common dollar baseline

A provider announcement or public page may establish a monthly price without establishing the final charge. Let P -posted denote that documented amount. A useful intermediate calculation is

$$IH_{jpt}^{posted} = \frac{P_{jpt}^{posted}}{W_{jt}} \quad (2)$$

Posted-price IH must retain its tax and eligibility flags. It is not interchangeable with checkout IH and is not automatically a lower bound if discounts, credits or other terms remain unknown. The documentary case in Section 6 uses this intermediate measure because the historical announcement does not resolve tax treatment.

For comparison with a regular US reference price B , let E be local currency units per US dollar. The common-price baseline is

$$IH_{jpt}^{USD} = \frac{B_{pt} E_{jt}}{W_{jt}} \quad (3)$$

This baseline asks how much work buys a common dollar amount at a reference exchange rate. It supports changes over time and a recognizable wage comparison. It does not estimate local purchasing conditions. The historical US\$20 illustration and the website's basket-v1 belong to this category.

For a matched plan and date, define a local-price factor L as the checkout charge relative to the converted US reference. Then

$$L_{jpt} = \frac{P_{jpt}}{B_{pt} E_{jt}} ; \quad IH_{jpt}^{local} = L_{jpt} IH_{jpt}^{USD} \quad (4)$$

L exposes the empirical price component absent from a common-price ranking. Depending on the specified tax basis, it can include taxes, regional prices and mandatory fees. It is not a pure measure of provider price discrimination. A separate tax-exclusive series would be needed to isolate that part of the comparison. Nor can L be interpreted as a matched-product price difference when the numerator and denominator refer to different tiers or bundles.

3.3 Matched service and paid entry cost

A matched-plan series compares the same named plan and documented version across markets. Its membership and restrictions are fixed for that comparison. A paid-entry series asks a different question: what is the least expensive qualifying paid subscription available to the specified purchaser? Let A be the set of eligible plans in an explicitly versioned catalog v . Then

$$IH_{jt}^{entry}(v) = \min_{p \in A_{jt}(v)} IH_{jpt}^{local} \quad (5)$$

Eligibility requires a regular consumer web subscription, recurring monthly billing, documented availability and a stated minimum service definition. Trials, student-only discounts, introductory prices, annual-payment equivalents, business contracts and API usage are separate comparisons. Advertising, quotas, language support, model access and bundled storage must be recorded; their presence does not imply equal practical value. The catalog should include qualifying lower-priced tiers, including ChatGPT Go and Google AI Plus where their local terms are verified. OpenAI's January 2026 Go announcement explicitly distinguishes its US price from localized prices (OpenAI, 2026).

The label is "lowest verified paid-entry cost within catalog v," not the universal cost of accessing AI. A free option is reported as free access rather than folded into a paid-only minimum. If eligible products are missing prices, the minimum among the observed subset is only the lowest observed offer. It is an upper bound on the minimum of a complete larger catalog under the same eligibility rules, not a verified market-wide minimum. Documented unavailability and an unverified offer are different states.

The two questions should be published alongside each other. Replacing a \$20 product with a cheaper, differently constrained product changes what is being purchased. A fixed-service comparison preserves continuity; the entry series records the changing cost of obtaining a qualifying paid option. Neither alone represents all dimensions of access.

4 Geographic scope and observation protocol

4.1 The wage jurisdiction is part of the result

A country is not necessarily a wage-setting unit. A capital-city floor must appear as a capital-city observation in the number's title, table row, map tooltip, shared image and export. "United States — Washington, DC — general non-tipped minimum wage" identifies a different population from "United States — federal statutory benchmark." An explanatory paragraph on a methods page cannot substitute for these labels at the point of use.

Where a federal floor coexists with higher applicable state or local floors, the federal amount is a legal reference rather than the prevailing wage for every covered worker. Tipped, apprentice, youth, sectoral and special employer rules require their own populations. The default denominator is the applicable general adult statutory floor, with an explicit age threshold and coverage statement. Collective agreements can be informative, but should not silently enter a ranking of general statutory minima.

National summaries require a defined target population. A range across verified jurisdictions can show geographic variation while disclosing incomplete coverage. An employment-weighted mean or median requires independently sourced employment weights and a decision about the workers being represented. Averaging jurisdictional wage rates and then dividing the price by that average is generally different from averaging workers' or jurisdictions' IH values. No national estimate should be manufactured from the capital value or from an unweighted midpoint.

An official hourly rate is used directly. Daily, weekly and monthly wages need evidenced hours-per-period conversions, including the treatment of paid rest and additional wage payments. A universal 174-hour monthly divisor is not justified. Statutory floors describe entitlements; violations, informal work and differences in actual earnings require separate evidence.

4.2 The minimum observation record

Each price record identifies the provider, plan version, market, eligible purchaser, billing channel, currency, recurring amount, tax status, unavoidable fees, availability and dated source. Each wage record identifies the jurisdiction, population, pay period, statutory basis, hourly conversion and

effective interval. Source capture, verification, editorial approval and legal or commercial effective dates are stored separately. A page retrieved in 2026 can document a 2023 launch; retrieval does not make that historical price a current offer.

Market identification cannot rely on language alone. A Korean-language URL may return a US offer, while an authenticated account may reveal prices unavailable to an anonymous visitor. A collector must preserve the market selector, displayed currency and conditions that establish eligibility. Public-page and checkout evidence should have distinct status fields. No historical final checkout is inferred from a later page or an unverified tax assumption.

Missing states remain explicit: no general statutory minimum; awaiting wage verification; collective reference; missing price; unresolved tax coverage; unsupported service; and availability not yet verified. Missing numerical inputs yield null, never zero. A verified free offer can have a zero subscription fee but belongs to the free-access series. It must not be confused with missing paid prices.

Price, wage and FX freshness policies should be published separately. The existing website uses review thresholds of 7 days for prices, 45 days for wages and 7 days for reference FX. These are operational review rules, not estimates of statistical uncertainty. A collection failure retains the previous value and its actual verification date. Local checkout observations will need an evidenced refresh policy and re-verification after relevant tax, eligibility or billing changes; this article does not claim that such a global collection service is already operating.

4.3 Versioning and comparability

A matched panel requires common reference dates, product definitions, tax treatment and wage populations. Membership changes create a new catalog or basket version. A plan redesign or age-eligibility change can be a series break even when a price or nominal wage remains the same. Both old and corrected observations should remain traceable. Daily history starts with an actual recorded observation, with no synthetic backfill presented as collection history.

A median across named plans is defensible as a specified reference purchase, not as evidence of product equivalence. If the price vector is common across markets, its median remains a common number. Adding lower-priced plans to that same common vector cannot create new cross-country price information. A local basket requires a complete matched set of verified local prices; an absent constituent does not silently change the basket.

5 Analytical implications and rank sensitivity

5.1 What a common-price ranking can identify

Write the dollar-converted hourly wage as $w = W/E$. For any common positive reference B , Equation 3 is B/w . Consequently, for two locations a and b , the ratio of their benchmark hours is $w-b$ divided by $w-a$, and the order is exactly reversed. This result holds for every common positive price and every common-price basket statistic. It does not depend on the size of the sample or on a regression fit.

A cross-sectional regression of log benchmark IH on log dollar wages therefore restates an accounting identity. The same logic prevents a change in the common benchmark from establishing a change in relative geographic affordability: it rescales all hours without changing the ranking. The empirical value of the benchmark is transparent wage purchasing power in an intelligible unit. Its rank is not new evidence about geographic subscription prices, demand, productivity or AI policy.

With verified local prices, both P and W can vary. Differences in local IH can then be described in terms of matched price and wage differences. This adds observed price information but does not, by itself, identify a causal effect of regional pricing. Taxes, market selection, bundles and eligibility must still be addressed. A complete local-price dataset is a prerequisite for testing hypotheses about these differences.

5.2 Changes over time and alternative denominators

For positive inputs with unchanged definitions, the local series has the exact accounting decomposition

$$\Delta \ln IH^{local} = \Delta \ln P - \Delta \ln W \quad (6)$$

Combining Equations 3 and 4 also gives the equivalent sum of changes in log B and log E, minus log W, plus log L. This expanded expression is a comparison with the US benchmark, not a causal model of exchange-rate pass-through. If a local nominal price is unchanged, a change in a reporting exchange rate does not mechanically change local IH. It changes the benchmark and L in offsetting directions.

Multiplying log differences by 100 gives log points. Ordinary percentage changes equal 100 times the exponential of the log difference minus one. Ratios calculated from annual-average inputs are not generally averages of monthly ratios. A frozen prior-year FX value creates a scenario, not an observed exchange-rate outcome.

A median-wage extension is valid only for matched periods, wage concepts and populations. With the Kaitz ratio defined as minimum divided by median hourly earnings, the identity is

$$IH_{jpt}^{median} = \kappa_{jt} IH_{jpt}^{minimum} \quad (7)$$

A national full-time earnings ratio cannot automatically convert a capital-jurisdiction minimum into a local median. The OECD indicator can inform a matched extension, but cannot repair missing local earnings data (OECD, n.d.-b). Common samples are also necessary before comparing rank changes.

Converting a local price and local wage by the same purchasing-power factor Q leaves the ratio unchanged:

$$\frac{\frac{P}{Q}}{\frac{W}{Q}} = \frac{P}{W} \quad (8)$$

Replacing only the market FX term in the benchmark with a PPP factor defines a counterfactual purchase. It does not recover an observed local price. Differences between regression slopes under these substitutions are not an identified causal decomposition of international price levels.

5.3 Small gaps and the precision of ranks

Exact ordering and substantive separation are different claims. Two stipulated exact values of 1.08 and 1.09 hours differ by 36 seconds. Raising the first benchmark's FX input alone by 0.926% would make their unrounded values equal; a 1% increase would reverse the order. This calculation is a sensitivity example using stipulated values, not evidence that any named country's observed difference is exchange-rate noise.

For a common-dollar ranking, report sensitivity to specified FX windows and relative perturbations, as well as the publication-date ordering. A 30-day range or alternative monthly average describes dependence on a stated window, not a sampling confidence interval. A local-price ranking needs sensitivities appropriate to local taxes, prices and wage coverage; adding artificial FX uncertainty to a same-currency ratio would be misleading.

When defensible positive bounds on a local price and wage are available, conservative bounds are

$$IH^{local} \in \left[\frac{P_{min}}{W_{max}}, \frac{P_{max}}{W_{min}} \right] \quad (9)$$

Location a is unambiguously below b under those bounds only if a's upper bound is below b's lower bound. Overlap means the order is not robust to the stated range, not that the locations are statistically identical. Correlated restrictions can narrow the feasible set; an arbitrary interval must not replace missing evidence. Publish absolute hours, differences in minutes, the definition of the bounds and coverage alongside any rank. Labels such as "most affordable country" exceed a capital-only, product-specific comparison.

6 Documentary illustration and reproducible baseline

6.1 Historical local posted prices and wage jurisdictions

Anthropic's 7 September 2023 announcement states monthly Claude Pro prices of US\$20 and UK£18. The US Department of Labor identifies the federal \$7.25 hourly floor and the precedence of higher applicable state floors. The District of Columbia announced a \$17.00 non-tipped rate effective 1 July 2023. The UK Low Pay Commission documents the £10.42 National Living Wage for ages 23 and over from 1 April 2023 (Anthropic, 2023; U.S. Department of Labor, n.d.; District of Columbia Department of Employment Services, 2023; Low Pay Commission, 2023).

These records support three documentary price-to-wage ratios at the launch date. The two US rows share a price and differ in the wage reference; they are not independent national observations. The UK and US wage populations are also not identical. Table 1 therefore reports clearly named observations rather than a country leaderboard.

Table 1. Claude Pro posted-price illustration at 7 September 2023

Wage reference	Posted monthly price	Gross wage per hour	Posted IH hours
US federal benchmark	USD 20.00	USD 7.25	2.759
Washington DC non-tipped	USD 20.00	USD 17.00	1.176
UK National Living Wage age 23+	GBP 18.00	GBP 10.42	1.727

Note. Historical documentary observations, retrieved 29 September 2026. The announcement does not establish tax-inclusive checkout charges. Checkout IH is therefore unavailable for all three rows. The federal benchmark must not replace a higher wage legally applicable to a particular US worker. Displayed ratios are rounded only after division.




Figure 1. The same US posted price produces different ratios under the federal benchmark and the District rate. Bars retain the order of Table 1 rather than ranking countries. The UK row uses its own posted currency amount and covered age group. These are posted-price ratios, not verified checkout costs or national averages.

The US federal ratio is 2.345 times the District ratio. Nothing about the subscription changed between those two calculations. Substituting the District observation for a US national headline would therefore conceal a consequential denominator choice. The UK ratio lies between the two US references, illustrating how a cross-country ordering can reverse solely through the geographic wage basis. This is a descriptive illustration of specification sensitivity, not an estimate of average American or British affordability.

No reference exchange rate is needed to divide each posted price by its same-currency wage. However, missing tax information prevents these observations from completing the principal checkout measure or identifying a tax-adjusted regional discount. This limitation is part of the result: the available documentary evidence supports Equation 2 and does not yet support Equation 1.

6.2 A separately reproduced common-price baseline

The Korean series pairs official hourly minimum wages with a frozen World Bank period-average exchange-rate response. At a fixed US\$20 reference, the resulting hours are shown in Table 2 (Minimum Wage Commission, n.d.; World Bank, n.d.). It demonstrates the accounting baseline and should not be described as a historical series of Korean checkout prices.

Table 2. Korean US\$20 reference-price baseline

Year	Gross wage KRW per hour	FX KRW per USD	Benchmark IH hours
2023	9,620	1,305.663	2.714
2024	9,860	1,363.375	2.765
2025	10,030	1,422.441	2.836

Note. Wages and FX inputs are retained at full precision in the research files. The table is rounded for display. Annual ratios use the annual-average FX input; they are not a monthly checkout-price panel.

From 2023 to 2025, the benchmark rises by 4.393 log points, equivalent to 4.491%. Exchange-rate depreciation contributes 8.566 log points and nominal wage growth offsets 4.174 log points. The decomposition is exact before rounding. It does not establish that Korean buyers' actual subscription expenditure increased by that amount. That conclusion would require local price observations and any relevant taxes or payment charges.

6.3 Lower-priced tiers and what rescaling does not show

At a fixed wage and FX input, substituting an \$8 reference for \$20 multiplies benchmark hours by 0.4 and leaves every geographic rank unchanged. The 60% reduction is algebra, not an estimated adoption effect. The \$8 example is grounded in OpenAI's US Go launch announcement, which also states that some markets have localized prices. It must not be transferred to all countries as their observed entry cost.

A verified local entry series could produce different results because prices, availability and eligibility vary. Establishing those differences requires the complete qualifying catalog for each covered market and date. Neither the present documentary illustration nor a website catalog of US prices establishes a global local-entry ranking. Dynamic plan pages, including Google AI plans, should be captured with their selected market, plan version and current terms rather than treated as timeless price constants (Google, n.d.).

6.4 The historical panel and the boundary of the evidence

The original materials cite OECD statutory minimum-wage data (OECD, n.d.-a) and contain a reported 37-country 2025 table with 36 nonmissing market-rate values, alongside regression, decomposition and rollout tables. Their underlying 111-row panel, full-precision inputs, conversion provenance and original estimation code were not supplied for this revision. Those tables remain preserved as source transcriptions, but have been removed from the main empirical argument and from the abstract's evidence claims.

A count or arithmetic check of rounded entries is not a reconstruction of the original study. Nor does private peer review substitute for independent replication. Claims about international elasticities, rollout sequencing and causal mechanisms require the original data and a justified estimation design.

This paper makes no new inferential claim from those tables and supplies no synthetic replacement panel.

7 Publication and research implementation

7.1 Reporting the correct object

A public result should identify the measure before the number: local checkout, posted price, common-dollar benchmark, or lowest verified paid entry. It should display the plan, jurisdiction, population, price basis, reference date and status adjacent to the result. Rank labels should name the comparison set and use unrounded values; exact ties may share a rank. Fine differences should be accompanied by their absolute magnitude and the relevant sensitivity analysis.

The existing website's operational series remains a US list-price benchmark unless and until verified local records are implemented. Its basket-v1 uses ChatGPT Plus, Claude Pro and Google AI Pro. This paper's local-price protocol is a research specification, not evidence that a new global local-price series has already replaced the website calculation. Initial launch dates also matter: the three subscriptions did not all exist in early 2023 (OpenAI, 2023; Anthropic, 2023; Hsiao, 2024).

A reasonable first local-price release is a limited set of markets with complete observation records, publication-date snapshots and explicit exclusions. Expansion should follow evidence coverage rather than a target country count. The missing-data register should distinguish absence of a service from failure to verify it. A complete catalog within a small defined scope is more interpretable than a broad ranking built from mixed price concepts.

7.2 Reproduction claims and public access to evidence

Reproduction has distinct levels: inspecting a formula; rerunning arithmetic from recorded inputs; tracing those inputs to archived primary evidence; and replicating an empirical panel and its estimates. This revision provides formulas, documentary inputs, explicit divisions and the Korean benchmark calculations. The original panel and regressions have not been reproduced. These levels should be named wherever reproducibility is claimed.

The public website currently distributes only the watermarked PDF. The editable source, input tables, source register and audit code are maintained with the research project; they are not a public downloadable replication package. Readers can recompute the printed documentary examples, but end-to-end computational replication from public files is not available in this release. A public repository or other accessible data-and-code deposit would be required before claiming that level of reproducibility. The phrase "Possible to reproduce" should not be used as an unqualified assurance covering unavailable panel data.

7.3 Policy interpretation and future tests

A future matched local-price panel can describe the variation in L, compare matched-plan and entry-cost distributions, and examine whether a common-dollar ranking changes under observed local

prices. Tests should specify the common sample, product versions, wage populations, tax treatment and reference dates before calculating rank differences. Missingness may be related to market access and must be analyzed rather than assumed random.

The resulting descriptive evidence could inform pricing or inclusion debates. It would not by itself identify welfare effects, optimal subsidies, labor displacement or the value of domestic compute investment. The Broadband Commission's broadband affordability target has its own service and income definition; it is not an AI subscription threshold (Broadband Commission, n.d.). IH does not impose an affordability pass/fail rule or treat the statutory wage as household disposable income.

8 Limitations and conclusion

The principal empirical limitation is the absence of a multi-market dataset of verified local checkout prices. The local estimand is defined, but the documentary example establishes only posted-price ratios. Historical tax coverage, payment charges and complete entry-plan catalogs remain unresolved. The Korean baseline is independently recalculated from documented inputs, while the original cross-country panel remains outside the verified empirical contribution.

Even with complete checkout records, statutory wages would not represent all workers, and plan labels would not guarantee equal capabilities. Wage enforcement, nonemployment, informal work, household needs, free access and the practical value of the service remain outside the ratio. Any national aggregation and any quality-adjusted extension require additional evidence and assumptions.

The revision narrows the claim and strengthens the research design. Intelligence Hours is a transparent way to express a defined subscription burden in working time. A common-price ranking is a rescaled wage ranking. Local checkout observations add a distinct price dimension, provided products, populations and dates are matched. The documentary case demonstrates the importance of jurisdiction labels and also the limits of evidence that stops at a posted price. The next empirical contribution is therefore a verified local-price panel with a complete stated catalog, not a larger leaderboard generated from a common dollar amount.

Declarations

Review and version history. The author confirms that the initial manuscript completed peer review in private circulation before public release. Version 1.1 responds to subsequent critical feedback on price identification, geographic scope, lower-priced plans and reproducibility. No journal referee process, journal acceptance or review of this revised version is represented as independently documented. Private peer review and journal publication are separate statuses.

Author contribution. Kwangseob Ahn developed the concept and supplied the initial research manuscript. The manuscript attributes conceptualization, methodology, data curation, formal analysis and writing to the author. Final responsibility for the text and claims rests with the named author.

Affiliation and interests. The author is an adjunct professor at Sejong University and the founder and principal consultant of INLEVEL9, which operates Intelligence Hours. The initial manuscript reports

no specific grant funding. Institutional affiliations do not imply endorsement by Sejong University or any cited organization.

Use of generative AI. Generative AI assisted with source discovery, restructuring, English-language editing, code preparation and arithmetic checks. AI tools are not authors. This public revision is not a journal submission; journal-specific disclosure and final submission approval remain the responsibility of the author.

Data and code availability. The public PDF contains the inputs and source links for the documentary example and Korean baseline. The project retains a three-row documentary input file, a frozen Korean FX response, an offline audit, source records and the six original table transcriptions. Those retained files are not downloadable from the public methods page. No public replication deposit DOI is claimed. Appendix A identifies the supplied materials and precisely limits what has been reproduced.

Author biography

Kwangseob Ahn is an adjunct professor in the School of Business at Sejong University in the Republic of Korea, where he teaches business data management, business analytics and statistical approaches to data-informed decision making. He has more than 15 years of experience in the IT industry and is the founder and principal consultant of INLEVEL9, a boutique consulting firm.

[INLEVEL9 author profile](#) · [LinkedIn](#)

References

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Ahia, O., Kumar, S., Gonen, H., Kasai, J., Mortensen, D. R., Smith, N. A., & Tsvetkov, Y. (2023). Do all languages cost the same? Tokenization in the era of commercial language models. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9904–9923. <https://doi.org/10.18653/v1/2023.emnlp-main.614>
- Anthropic. (2023, September 7). Introducing Claude Pro. <https://www.anthropic.com/news/claude-pro>
- Ashenfelter, O. (2012). Comparing real wage rates: Presidential address. *American Economic Review*, 102(2), 617–642. <https://doi.org/10.1257/aer.102.2.617>
- Barrantes, R., & Galperin, H. (2008). Can the poor afford mobile telephony? Evidence from Latin America. *Telecommunications Policy*, 32(8), 521–530. <https://doi.org/10.1016/j.telpol.2008.06.002>
- Broadband Commission for Sustainable Development. (n.d.). Advocacy target 2: Make broadband affordable. Retrieved September 29, 2026, from <https://www.broadbandcommission.org/advocacy-targets/2-affordability/>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Cavallo, A., & Rigobon, R. (2016). The Billion Prices Project: Using online prices for measurement and research. *Journal of Economic Perspectives*, 30(2), 151–178. <https://doi.org/10.1257/jep.30.2.151>
- Deneckere, R. J., & McAfee, R. P. (1996). Damaged goods. *Journal of Economics & Management Strategy*, 5(2), 149–174. <https://doi.org/10.1111/j.1430-9134.1996.00149.x>
- District of Columbia Department of Employment Services. (2023, June 29). Mayor Bowser announces minimum wage increase effective July 1, 2023. <https://does.dc.gov/release/mayor-bowser-announces-minimum-wage-increase-effective-july-1-2023>
- Google. (n.d.). Google AI plans. Retrieved September 29, 2026. https://one.google.com/intl/en_us/about/google-ai-plans/
- Hilbert, M. (2010). When is cheap, cheap enough to bridge the digital divide? Modeling income related structural challenges of technology diffusion in Latin America. *World Development*, 38(5), 756–770. <https://doi.org/10.1016/j.worlddev.2009.11.019>
- Hsiao, S. (2024, February 8). Bard becomes Gemini: Try Ultra 1.0 and a new mobile app today. Google. <https://blog.google/products-and-platforms/products/gemini/bard-gemini-advanced-app/>
- Humlum, A., & Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1), e2414972121. <https://doi.org/10.1073/pnas.2414972121>
- Kanbur, R., & Ronconi, L. (2018). Enforcement matters: The effective regulation of labour. *International Labour Review*, 157(3), 331–356. <https://doi.org/10.1111/ilr.12112>
- Low Pay Commission. (2023, March 31). The National Minimum Wage in 2023. <https://www.gov.uk/government/publications/the-national-minimum-wage-in-2023/the-national-minimum-wage-in-2023>
- Minimum Wage Commission. (n.d.). Annual minimum wage decisions [Data table, Republic of Korea]. Retrieved September 29, 2026, from <https://www.minimumwage.go.kr/minWage/policy/decisionMain.do>

- OECD. (n.d.-a). Statutory minimum wages in national currency [Data set, DSD_EARNINGS@MW_CURP]. https://sdmx.oecd.org/public/rest/v1/data/OECD.ELS.SAE,DSD_EARNINGS@MW_CURP,1.0/
- OECD. (n.d.-b). Minimum wages relative to median wages [Indicator; source manuscript identifies DSD_EARNINGS@MIN2AVE]. <https://doi.org/10.1787/c9e94817-en>
- OpenAI. (2023, February 1). Introducing ChatGPT Plus. <https://openai.com/index/chatgpt-plus/>
- OpenAI. (2026, January 16). Introducing ChatGPT Go, now available worldwide. <https://openai.com/index/introducing-chatgpt-go/>
- Parsley, D. C., & Wei, S.-J. (2007). A prism into the PPP puzzles: The micro-foundations of Big Mac real exchange rates. *The Economic Journal*, 117(523), 1336–1356. <https://doi.org/10.1111/j.1468-0297.2007.02085.x>
- Stone, M. E. (2006). What is housing affordability? The case for the residual income approach. *Housing Policy Debate*, 17(1), 151–184. <https://doi.org/10.1080/10511482.2006.9521564>
- U.S. Department of Labor. (n.d.). Fair Labor Standards Act Advisor: What is the minimum wage? Retrieved September 29, 2026. <https://webapps.dol.gov/elaws/faq/esa/flsa/001.htm>
- World Bank. (n.d.). Official exchange rate, LCU per US dollar, period average [World Development Indicators, PA.NUS.FCRF; Korea 2023–2025]. Retrieved September 29, 2026, from <https://api.worldbank.org/v2/country/KOR/indicator/PA.NUS.FCRF?date=2023:2025&format=json>
- Zeira, J. (1998). Workers, machines, and economic growth. *The Quarterly Journal of Economics*, 113(4), 1091–1117. <https://doi.org/10.1162/003355398555847>

Appendix A Evidence and reproducibility register

Table A1. Evidence supported by this release

Material	Observation scope	Check performed	Remaining limit
Local documentary case	Two launch prices and three wage references at 7 September 2023	Primary pages checked; three divisions rerun	Tax coverage unresolved; no checkout IH
Korean benchmark	2023–2025 official hourly wages and frozen annual FX	Inputs retained; levels and log contributions rerun	Common USD scenario, not local transaction prices
Rank sensitivity	Stipulated 1.08 and 1.09 hours	Algebra and perturbation calculation	Synthetic example, not observed uncertainty
Original cross-country tables	Six author-supplied summary tables	Transcription and rounding checks	Original panel and estimates not reproduced

The original files supplied in Markdown, Word and PDF remain identified by their SHA-256 digests. The three documentary rows are stored in `local-price-case.csv`; missing checkout amounts are blank and the audit returns null for checkout hours. The independent Korean inputs are stored in `korea-inputs.csv` and `worldbank-korea-fx.json`. Historical source transcriptions retain their original descending ranks and are not inputs to the revised article's local-price results.

The internal audit runs with Python 3 using `python3 verify-research.py --data-dir data`. It recalculates the documentary case and the Korean baseline, checks the exact decomposition and verifies the common-price ordering, local-factor identity, Kaitz identity and PPP cancellation on explicitly synthetic inputs. It also demonstrates the 36-second gap and 0.926% relative perturbation example. Tests of algebra are not observations. The audit does not collect missing checkout prices or recreate the original regression and rollout datasets.

The historical Korean 2026 row retained in the source files uses the official KRW10,320 wage with 2025 FX carried forward. It is a scenario, excluded from the observed 2023–2025 series in Table 2. Re-running the audit does not refresh input vintages or turn a scenario into an observed year.