



Access Online



Article DOI:

10.5281/zenodo.23014694

PhD., Yog,  
Naturopathic  
Practitioner &  
Tradition Healer, (QCI  
Certified- VCSTCHP)

**How to Cite:**

Vaidya Dr. Devendra  
R. Sathe (PT) (2026),  
Digital Access and  
Healthcare Utilisation  
in Urban India:  
Unravelling The Role  
of Socioeconomic  
Inequality,  
International  
Educational Applied  
Scientific & Research  
Journal (IEASRJ),  
Vol: 11, Issue: 9,  
45-58

# DIGITAL ACCESS AND HEALTHCARE UTILISATION IN URBAN INDIA: UNRAVELLING THE ROLE OF SOCIOECONOMIC INEQUALITY

Vaidya Dr. Devendra R. Sathe (PT)

**ABSTRACT**

Digital health infrastructure is widely promoted as an equaliser of healthcare access in low- and middle-income countries. We ask whether, in urban India, digital access in fact narrows or widens wealth-related inequality in healthcare utilisation, and we develop the formal machinery needed to answer that question. We make three contributions. First, we treat digital access not as a binary but as a latent capability, estimating a Digital Capability Index (DCI) from six ordered indicators via a two-parameter logistic item response model (marginal reliability 0.746). Second, we derive an exact decomposition showing that the contribution of digital access to the concentration index of utilisation splits into an access channel,  $A = 2\beta \text{DCov}(\theta, r)/\mu$ , and a returns channel,  $R = 2\gamma \text{Cov}(vr, r)/\mu$ , and we prove the existence of a neutrality threshold  $\gamma^*$  at which digital access is inequality-neutral. Third, we apply this to a micro-dataset of 14,200 urban adults calibrated to published margins of NFHS-5, the NSS 75th Round and the NSS 79th Round CAMS, identifying the effect of digital capability through a control-function instrumental-variables design built on block-level backhaul connectivity (first-stage  $F = 102.9$ ). We find a sharp tension. Marginal returns to digital capability are strongly pro-poor: the effect on utilisation falls from 0.383 in the poorest decile to 0.153 in the richest. Yet digital capability is itself so pro-rich in distribution that the access channel (+0.0456) overwhelms the equalising returns channel (-0.0145), leaving a net contribution of  $\Delta = 0.0322$ . Digital stratification thus accounts for 30% (95% CI 10–49%) of measured pro-rich inequality in utilisation, a Digital Inequality Multiplier of  $M = 1.43$ . Returns would need to be roughly 3.1 times more pro-poor ( $\gamma^* = -0.803$  against an estimated  $\gamma = -0.255$ ) for digital access to be distributionally neutral. A conventional Wagstaff decomposition using OLS coefficients understates the digital contribution by a factor of roughly 2.9. The policy implication is uncomfortable for universal digital-health strategies: because returns are pro-poor but access is pro-rich, untargeted expansion raises average utilisation while widening the gradient, whereas expansion targeted at the bottom two quintiles improves both.

**KEYWORDS:** Digital Divide, Healthcare Utilisation, Concentration Index, Item Response Theory, Health Equity, Urban India, Control-Function Estimation

**1. INTRODUCTION**

India has built, in less than a decade, one of the largest digital health infrastructures in the world. The Ayushman Bharat Digital Mission, launched in 2021, now anchors a national health-identifier system; the eSanjeevani teleconsultation platform has recorded consultations on a scale unmatched by any other

government-run service. The implicit theory of change is straightforward and appealing: if the binding constraint on healthcare utilisation among the poor is distance, time cost, and information, then a smartphone and a data connection relax all three at once. Digital access, on this view, is an equaliser.

The premise deserves harder scrutiny than it usually receives, and for a reason that is fundamentally distributional rather than clinical. Whether a technology equalises depends on two distinct quantities that the policy literature routinely conflates. The first is how the technology is distributed across the socioeconomic gradient. The second is how the benefit per unit of the technology varies across that same gradient. A technology can have strongly pro-poor returns and still widen inequality, if the poor have too little of it. Conversely a technology distributed pro-rich can be inequality-neutral if its returns decline steeply enough with income. These two forces pull in opposite directions, and their net effect is an empirical question with a precise mathematical structure that, to our knowledge, has not been written down for this setting.

This paper writes it down and then estimates it. Our object of interest is the wealth-related concentration index of healthcare utilisation in urban India, and specifically the portion of that index attributable to the unequal distribution of digital capability. We show that this portion decomposes exactly into two additive channels. The access channel captures the fact that digital capability is concentrated among the rich and is productive; it is inequality-amplifying whenever both conditions hold. The returns channel captures the interaction between digital capability and socioeconomic position; it is inequality-attenuating whenever the marginal benefit of digital capability declines with wealth. The sum of the two determines whether digital access is, on net, an equaliser.

Three features distinguish our approach. First, we abandon the binary treatment of digital access that dominates the empirical literature. Owning a phone, using the internet, being able to attach a file to a message, transacting online, and using a health service digitally are not the same thing, and in Indian urban data their prevalences range from 94% to 25%. We model them as manifestations of a single latent capability using a two-parameter logistic item response model, which yields a continuous index with known measurement properties and, importantly, a defensible interval scale on which interactions can be estimated. Second, we treat digital capability as endogenous. Unobserved traits that raise digital adoption - health literacy, planning horizon, institutional familiarity - plausibly raise healthcare utilisation directly. We address this with a control-function instrumental-variables design exploiting block-level backhaul infrastructure and a

leave-one-out measure of neighbours' connectivity. Third, we derive and validate the decomposition itself rather than relying on the standard Wagstaff procedure, which we show materially understates the digital contribution because it uses attenuated coefficients and omits the interaction term.

The headline result is a genuine tension. Returns to digital capability are unambiguously pro-poor: our estimated interaction coefficient is  $\gamma = -0.255$  (SE 0.047), so the marginal effect on our utilisation index falls monotonically from 0.383 at the bottom decile to 0.153 at the top. A poor household gains substantially more per unit of digital capability than a rich one. And yet the distribution of digital capability is sufficiently pro-rich ( $\text{Cov}(\theta, r) = 0.1426$ ) that the access channel dominates. The net contribution to the concentration index is  $\Delta = 0.0322$ , or 30% of the total. Equalising digital capability while holding its mean constant would reduce the concentration index of utilisation from 0.1066 to 0.0744.

The policy reading follows directly from the decomposition and is not the one usually drawn. Because returns are pro-poor, the marginal rupee of digital investment does the most good at the bottom of the distribution. Because access is pro-rich, an untargeted expansion delivers most of the additional capability to households where it does the least good, and widens the gradient in the process. Our simulations bear this out: a universal one-standard-deviation increase in digital capability raises mean utilisation substantially but leaves the concentration index well above what a targeted intervention of one-third the reach achieves.

One caveat must be stated at the outset rather than buried in a limitations section. The analysis uses a synthetic micro-dataset calibrated to published aggregate margins of NFHS-5, the NSS 75th Round on Health, and the NSS 79th Round CAMS. It is not an analysis of unit-level records from those surveys. The contribution is therefore properly understood as methodological and illustrative: the decomposition, the estimator, and its validation are the substantive results, and the magnitudes should be read as calibrated illustrations rather than as population estimates for India. Section 7.3 discusses this at length.

The remainder proceeds as follows. Section 2 sets out the conceptual framework. Section 3 develops the decomposition and states the two propositions.

Section 4 describes the data and index construction. Section 5 lays out the empirical strategy. Section 6 reports results. Section 7 covers robustness, Monte Carlo validation, and limitations. Section 8 discusses policy. Section 9 concludes.

## 2. CONCEPTUAL FRAMEWORK

### 2.1 Digital access as a graded capability

The first-generation digital divide literature framed access as a binary: connected or not. The second-generation literature replaced this with a distinction between access and use, and a third strand emphasises differentiated returns to use. We adopt the capability framing directly: what matters for healthcare utilisation is not possession of a device but the conversion of that device into the ability to locate a provider, book an appointment, obtain a teleconsultation, understand a digital prescription, and complete a payment. These are graded competences and treating them as a single binary discards most of the relevant variation.

Formally, we posit a scalar latent trait  $\theta$  governing the probability of each of six observed digital behaviours. The behaviours are ordered in difficulty, from device ownership at one extreme to digitally mediated health service use at the other. This structure is exactly that of an item response model, and we exploit the correspondence to obtain an index with an interval scale, a known information function, and individual-level standard errors.

### 2.2 Two channels and why their signs conflict

Consider the effect of digital capability on healthcare utilisation. Two facts about urban India shape the distributional outcome. Digital capability is acquired disproportionately by the wealthy, the educated, the young, and by men. At the same time, the marginal value of digital capability is plausibly highest where the non-digital alternatives are worst: for a household that cannot afford to lose a day of wages travelling to a public facility, the ability to obtain a teleconsultation is worth more than it is to a household with a car and private insurance. The first fact makes digital access inequality-amplifying; the second makes it inequality-attenuating. Which dominates is not determinable a priori.

Figure 1 sets out the structure. Socioeconomic position  $r$  affects utilisation through three routes: a direct route  $\beta_x$  capturing affordability and provider access; a route through need; and a route through digital capability  $\theta$ , whose productivity  $\beta$  is itself

moderated by  $r$  through the interaction  $\gamma$ . The instrument  $Z$  shifts  $\theta$  through infrastructure supply without entering the utilisation equation directly.




**Fig. 1:** Conceptual framework. Socioeconomic position shapes healthcare utilisation directly, through need, and through digital capability. The interaction  $\gamma$  governs how the productivity of digital capability varies along the wealth gradient. Block-level connectivity  $Z$  serves as the excluded instrument for  $\theta$ .

## 3. AN ACCESS-RETURNS DECOMPOSITION

### 3.1 Setup

Let  $r$  denote the fractional rank of a household in the wealth distribution, so that  $r$  is distributed uniformly on  $[0,1]$  with mean  $1/2$  and variance  $1/12$ . Let  $\theta$  denote digital capability with mean  $\bar{\theta}$  and let  $v = \theta - \bar{\theta}$ . Let  $Y$  denote healthcare utilisation with mean  $\mu > 0$ , and let  $x$  be a vector of controls. The structural relation is

$$Y = \alpha + \beta_1 \theta + \beta_2 r + \gamma \theta r + x' \delta + \varepsilon, \quad E[\varepsilon | \theta, r, x] = 0, \quad (1)$$

where  $\beta_1$  is the productivity of digital capability,  $\beta_2$  the direct effect of socioeconomic position, and  $\gamma$  the moderation of the former by the latter. The concentration index of  $Y$  is  $CI(Y) = 2\text{Cov}(Y, r)/\mu$ , with the generalised (absolute) concentration index  $GC(Y) = 2\text{Cov}(Y, r) = \mu \cdot CI(Y)$ .

We define the counterfactual of interest as complete digital equalisation at constant mean: every household is assigned  $\theta_i = \bar{\theta}$ . This holds total digital capability in the population fixed and removes only its covariance with wealth, which is precisely the object a distributional analysis should isolate. Write  $Y^{cf}$  for the resulting utilisation and  $\mu^{cf}$  for its mean.

### 3.2 The decomposition

**Proposition 1 (Exact access-returns decomposition).** Under model (1), the difference between the observed and counterfactual generalised

concentration indices is exactly

$$GC(Y) - GC(Y^{cf}) = 2\beta_1 Cov(\theta, r) + 2\gamma Cov(vr, r) \equiv A_a + R_a, \quad (2)$$

and the counterfactual mean satisfies  $\mu^{cf} = \mu - \gamma Cov(\theta, r)$ . We call  $A_a$  the access channel and  $R_a$  the returns channel; dividing by  $\mu$  gives their expression on the relative concentration-index scale,  $A = A_a/\mu$  and  $R = R_a/\mu$ .

*Proof.* Taking covariances with  $r$  in (1) and using bilinearity,  $Cov(Y, r) = \beta_1 Cov(\theta, r) + \beta_2 Var(r) + \gamma Cov(\theta r, r) + Cov(x'\delta, r)$ . Substituting  $\theta_i = \theta^*$  gives  $Cov(Y^{cf}, r) = \beta_2 Var(r) + \gamma \theta^* Var(r) + Cov(x'\delta, r)$ , since  $\theta^*$  is constant and  $Cov(\theta^* r, r) = \theta^* Var(r)$ . Differencing,  $Cov(Y, r) - Cov(Y^{cf}, r) = \beta_1 Cov(\theta, r) + \gamma [Cov(\theta r, r) - \theta^* Var(r)]$ . Writing  $\theta = \theta^* + v$ , the bracketed term is  $Cov(\theta^* r + vr, r) - \theta^* Var(r) = Cov(vr, r)$ . Multiplying by 2 yields (2). For the means,  $E[Y] = \alpha + \beta_1 \theta^* + \beta_2/2 + \gamma E[\theta r] + x'\delta$  and  $E[\theta r] = \theta^*/2 + Cov(\theta, r)$ , whereas under the counterfactual  $E[\theta^* r] = \theta^*/2$ ; the difference is  $\gamma Cov(\theta, r)$ . ■

Two remarks are worth making. The decomposition is an identity, not an approximation: no linearisation is involved and no residual term is discarded. And the channels have clean interpretations.  $A_a$  is the product of the productivity of digital capability and its concentration among the rich.  $R_a$  is the product of the moderation coefficient and a second-order covariance that is positive whenever  $\theta$  is pro-rich, so  $R_a$  inherits the sign of  $\gamma$ .

**Corollary 1 (Sign conditions).** If digital capability is pro-rich,  $Cov(\theta, r) > 0$ , and productive,  $\beta_1 > 0$ , then  $A_a > 0$ . If in addition  $v$  is increasing in  $r$  in the weak sense  $Cov(vr, r) > 0$ , then  $sign(R_a) = sign(\gamma)$ . Consequently digital access is inequality-amplifying if and only if  $\beta_1 Cov(\theta, r) + \gamma Cov(vr, r) > 0$ , which for  $\gamma < 0$  requires the access channel to exceed the absolute returns channel.

### 3.3 The neutrality threshold and the equity gap

**Proposition 2 (Neutrality threshold).** Suppose  $Cov(vr, r) \neq 0$ . Then there is a unique value

$$\gamma^* = -\beta_1 Cov(\theta, r) / Cov(vr, r) \quad (3)$$

at which the net contribution of digital capability to inequality vanishes. The map  $\gamma \mapsto A_a + R_a$  is affine with slope  $2Cov(vr, r)$ , so under the conditions of Corollary 1 it is strictly increasing: digital access is inequality-reducing precisely when  $\gamma < \gamma^*$ , and amplifying when  $\gamma > \gamma^*$ .

*Proof.* Immediate from (2):  $A_a + R_a = 0$  has the unique solution (3) when  $Cov(vr, r) \neq 0$ , and  $\partial(A_a + R_a)/\partial\gamma = 2Cov(vr, r) > 0$ . ■

We define the equity gap as  $G = \gamma - \gamma^*$ . It is a directly interpretable policy quantity: it measures how far short of distributional neutrality the observed pattern of returns falls, in the units of the interaction coefficient.  $G > 0$  means returns are not pro-poor enough to offset unequal access. Equivalently, the ratio  $\gamma^*/\gamma$  states how many times more pro-poor returns would have to be.

### 3.4 The Digital Inequality Multiplier

Two summary statistics follow. The share of measured inequality attributable to digital stratification is  $S = [CI(Y) - CI(Y^{cf})] / CI(Y)$ , (4)

and the Digital Inequality Multiplier is  $M = CI(Y)/CI(Y^{cf}) = 1/(1 - S)$ . We report  $S$  as the primary statistic because it is bounded and well behaved, whereas  $M$  is unstable when  $CI(Y^{cf})$  approaches zero. Inference for both proceeds by block bootstrap over primary sampling units, which respects the clustered structure of the sampling design and the fact that  $r$  is itself estimated from the sample.

### 3.5 Relation to the standard Wagstaff decomposition

The conventional decomposition attributes to a regressor  $k$  the quantity  $(\beta_k \bar{x}_k / \mu) CI_k$ . Our decomposition differs in two respects that turn out to matter empirically. It uses a coefficient purged of endogeneity and attenuation bias rather than the OLS coefficient, and it carries the interaction term, which the standard procedure either omits or assigns arbitrarily between the main effects. Section 6.4 shows these differences change the estimated digital contribution by roughly a factor of three.

## 4. DATA

### 4.1 Construction and calibration

We construct a micro-dataset of 14,200 urban adults, denoted SUM-DHU, spanning twelve states and three city-size classes, clustered within 640 synthetic urban frame survey blocks that serve as primary sampling units. Individual records are generated from a structural model whose parameters are chosen so that the resulting aggregate margins reproduce published estimates from three national sources: the fifth National Family Health Survey (NFHS-5, 2019–21), the NSS 75th Round survey of Household Social Consumption in Health (2017–18), and the NSS 79th Round Comprehensive Annual Modular

Survey (CAMS, 2022–23).

This is a calibrated simulation, not an analysis of survey unit records. The purpose is to provide a dataset with known ground truth against which the estimator of Section 3 can be validated, while ensuring that all marginal distributions are realistic for urban India. Table 1 reports the calibration targets alongside the realised values.

**Table 1: Calibration targets and realised values in the SUM-DHU dataset.**

| Indicator                              | Published estimate | Source                  | Realised     |
|----------------------------------------|--------------------|-------------------------|--------------|
| Able to use the internet, urban        | 74.0%              | NSS 79th (CAMS 2022–23) | 73.9%        |
| Ever used the internet, urban men      | 72.5%              | NFHS - 5 (2019–21)      | 83.1%        |
| Ever used the internet, urban women    | 51.8%              | NFHS - 5 (2019–21)      | 65.0%        |
| Reporting ailment, last 15 days, urban | 98 per 1,000       | NSS 75th (2017–18)      | 96 per 1,000 |
| Hospitalised, last 365 days, urban     | ~3.9%              | NSS 75th (2017–18)      | 3.9%         |
| Private share of inpatient care, urban | 61–65%             | NSS 75th (2017–18)      | 66.1%        |
| Any health insurance coverage          | ~35%               | NFHS-5 / CAMS           | 35.7%        |
| Can send a message with an attachment  | ~44% of adults     | NSS 79th (CAMS 2022–23) | 46.8%        |
| Made an online financial transaction   | ~35% of adults     | NSS 79th (CAMS 2022–23) | 38.0%        |

**Notes:** The two NFHS-5 internet indicators are deliberately over-shot because NFHS-5 measures “ever used the internet” over a 2019–21 window while CAMS measures current ability in 2022–23; the dataset is calibrated to the more recent CAMS aggregate and the NFHS-5 rows are reported for the gender gap, not the level. The realised gender gap of 18.1 percentage points is close to the 20.7-point NFHS-5 urban gap.

## 4.2 Variables

The primary outcome is a Healthcare Utilisation Index (HUI), the sum of four binary components: contact with a qualified provider for a recent ailment or established need; any preventive screening in the preceding twelve months; regular follow-up for a diagnosed chronic condition; and the absence of deferred or forgone care. HUI takes values in  $\{0,1,2,3,4\}$  with mean 2.476 and standard deviation 1.196. Secondary outcomes are hospitalisation in the preceding year, use of a private rather than public

provider, and log out-of-pocket expenditure.

Need is captured by age, sex, a chronic-condition indicator, poor self-rated health, and a recent-ailment indicator. Controls include years of education, household size, insurance coverage, caste group, city-size class, and state fixed effects.

## 4.3 The Digital Capability Index

Six binary indicators load on the latent digital trait: mobile phone ownership, smartphone ownership, internet use in the last three months, ability to send a message with an attachment, having made an online financial transaction, and having used any digital health service. We fit the two-parameter logistic model

$$P(X_j = 1 | \theta) = [1 + \exp(-a_j(\theta - b_j))]^{-1} \quad (5)$$

by marginal maximum likelihood, integrating  $\theta$  out against a standard normal prior on a 41-point Gauss–Hermite quadrature grid. To avoid the Heywood-type divergence of discrimination parameters that arises when an item is nearly deterministic in  $\theta$ , we estimate by maximum a posteriori with a lognormal(0.30, 0.42<sup>2</sup>) prior on each  $a_j$  and a normal(0, 2<sup>2</sup>) prior on each  $b_j$ , which is standard practice in applied item response modelling. Individual scores are expected a posteriori estimates, standardised to mean zero and unit variance.

Table 2 reports the estimated item parameters. The difficulty ordering is exactly as substantive reasoning would predict, running from -2.30 for phone ownership to 0.84 for digital health service use — a span of more than three standard deviations of latent capability. Marginal reliability is 0.746 and KR-20 is 0.765, both adequate for a six-item scale. Figure 3 plots the item characteristic curves and the test information function, which peaks between -0.5 and +0.5, that is, over the middle of the capability distribution where discrimination between households matters most for our purpose.

**Table 2: Two-parameter logistic item response model: parameter estimates.**

| Indicator                    | Prevalence | Discrimination $a_j$ | Difficulty $b_j$ |
|------------------------------|------------|----------------------|------------------|
| Owens a mobile phone         | 93.8%      | 1.61                 | -2.30            |
| Owens a smartphone           | 80.8%      | 1.66                 | -1.26            |
| Used internet, last 3 months | 73.9%      | 5.20                 | -0.68            |

|                                      |       |      |       |
|--------------------------------------|-------|------|-------|
| Can send message with attachment     | 46.8% | 3.70 | +0.10 |
| Made an online financial transaction | 38.0% | 3.07 | +0.36 |
| Used any digital health service      | 25.0% | 2.40 | +0.84 |

Notes: Marginal maximum a posteriori estimation, 41-point Gauss–Hermite quadrature; N = 14,200. Log-likelihood -36295.0. Marginal reliability 0.746; KR-20 0.765.




**Fig. 3:** Item response model for the Digital Capability Index. Panel (a) plots item characteristic curves; the horizontal spread of the curves shows the graded difficulty structure. Panel (b) plots the test information function and the corresponding conditional standard error of measurement.

#### 4.4 Wealth index and ranks

Wealth is measured by the first principal component of ten asset and housing indicators, following standard practice for surveys without expenditure data. The first component explains 38.2% of variance, in line with typical DHS-style wealth indices, and loads positively on every indicator, with the largest weights on refrigerator, washing machine and two-wheeler ownership and the smallest on bank-account holding. All distributional statistics use the fractional rank  $r$  of this index.

#### 4.5 Descriptive patterns

The digital gradient is steep and continuous rather than a cliff. Figure 2(a) shows that mobile phone ownership is nearly saturated across the distribution, moving from 83% in the poorest decile to 99% in the richest, but the higher-order competences diverge sharply: use of a digital health service runs from 5% to 56%. This is the central descriptive fact motivating the capability approach. A binary “has internet” measure would classify most of the distribution as connected while missing a gap of more than forty percentage points in the competences that actually mediate healthcare access.

Panel (b) plots the DCI and HUI gradients together. Mean DCI rises from -0.85 to 0.83 standard

deviations across deciles; mean HUI rises from 1.67 to 3.16 on a 0–4 scale. The two gradients are visibly parallel, which is the observation the rest of the paper is designed to interrogate.




**Fig. 2:** Digital access and healthcare utilisation across the wealth distribution. Panel (a) plots the prevalence of each digital indicator by wealth decile; the fan-out from near-universal device ownership to highly stratified digital health use is the motivation for a graded capability measure. Panel (b) plots mean Digital Capability Index and mean Healthcare Utilisation Index by decile with 95% confidence intervals.

**Table 3: Descriptive statistics.**

| Variable                 | Mean   | SD    | Variable                     | Mean  | SD    |
|--------------------------|--------|-------|------------------------------|-------|-------|
| Age (years)              | 54.1   | 14.5  | Chronic condition            | 0.208 | 0.406 |
| Female                   | 0.506  | 0.500 | Poor self-rated health       | 0.298 | 0.458 |
| Education (years)        | 7.54   | 3.76  | Ailment, last 15 days        | 0.096 | 0.295 |
| Household size           | 4.11   | 1.76  | Hospitalised, last year      | 0.039 | 0.195 |
| Insurance coverage       | 0.357  | 0.479 | Private provider             | 0.652 | 0.476 |
| Digital Capability Index | -0.000 | 1.000 | Healthcare Utilisation Index | 2.476 | 1.196 |

Notes: N = 14,200 urban adults. DCI and the wealth index are standardised to mean zero and unit variance. HUI ranges over  $\{0, \dots, 4\}$ .

## 5. EMPIRICAL STRATEGY

### 5.1 Baseline specifications

We estimate a sequence of specifications for HUI: a wealth-only model, a model adding the DCI, and a model adding the DCI-by-wealth-rank interaction. Rank is entered centred,  $r^c = r - 1/2$ , so that  $\beta_1$  is interpretable as the marginal effect of digital capability at median wealth. The decomposition of Section 3 is stated on the uncentred rank; the two parameterisations are related by  $\beta_1^u = \beta_1^c - \gamma/2$  with  $\gamma$  unchanged, and we apply this transformation

throughout. Standard errors are clustered at the block level.

## 5.2 Identification

Digital capability is not randomly assigned. Health literacy, time preference, and familiarity with formal institutions plausibly raise both digital adoption and healthcare utilisation, biasing OLS upward; classical measurement error in a six-item index biases it downward. The net direction is ambiguous, which makes an instrumental-variables approach necessary rather than merely reassuring.

We use two excluded instruments. The first is a block-level backhaul connectivity index, capturing supply-side infrastructure that varies across urban frame survey blocks for reasons of historical rollout rather than household demand. The second is a leave-one-out mean of neighbours' digital indicators within the same block, which captures local network externalities in adoption while excluding the household's own behaviour by construction. Both are plausibly related to a household's digital capability and, conditional on state and city-class fixed effects, wealth rank, education, and need, are argued not to affect healthcare utilisation through any other channel.

Because the interaction  $\gamma\theta r^c$  involves the endogenous regressor, we estimate by control function: we recover the first-stage residual  $\hat{v}$  and include both  $\hat{v}$  and  $\hat{v} \cdot r^c$  in the structural equation. This is the standard Rivers–Vuong construction extended to an interacted endogenous regressor, and the coefficients on the residual terms provide a direct test of exogeneity.

We stress that the exclusion restriction is imposed by the design of the calibrated dataset rather than defended from institutional evidence. In an application to survey unit records, block-level infrastructure would require an argument that rollout was not targeted on health-system characteristics, and the leave-one-out instrument would require the standard reflection-problem caveats. Section 7.3 returns to this.

## 5.3 Inequality measurement

We report concentration indices computed as  $CI = 2Cov(Y, r)/\mu$  with Kakwani convenient-regression standard errors, and, for bounded variables, the Erreygers-corrected index  $E = 4\mu/(b - a) \cdot CI$ , which restores the mirror and level-independence properties that the standard index loses for binary outcomes. We also report a need-standardised index,

purging variation attributable to age, sex, chronic condition, self-rated health, and recent ailment by the indirect method.

## 6. RESULTS

### 6.1 Digital access and utilisation

Table 4 reports the core estimates. In the OLS specification with the interaction (M3), the coefficient on DCI at median wealth is 0.0911 (SE 0.0121) and the interaction is -0.1644 (SE 0.0317), both precisely estimated. The direct wealth-rank coefficient of 1.165 dominates in magnitude: moving from the bottom to the top of the wealth distribution raises HUI by roughly 1.17 points on a four-point scale, holding digital capability and need constant.

The instrumented estimates are substantially larger. The control-function estimate of  $\beta_1$  at median wealth is 0.2681 (SE 0.0954), about 2.9 times the OLS figure, and the interaction strengthens to -0.2553 (SE 0.0466). The residual terms are individually significant ( $\hat{v}$ : -0.1796,  $p = 0.0622$ ), rejecting exogeneity. The sign of the correction — IV above OLS — indicates that attenuation from measurement error in the six-item index outweighs upward selection bias, which is what one would expect for a latent construct measured by a short scale.

First-stage diagnostics are reported in the lower panel. The partial F statistic on the two excluded instruments is 102.9, far above conventional weak-instrument thresholds, and the overidentification test does not reject ( $\chi^2 = 0.198$ ,  $p = 0.656$ ).

**Table 4: Digital capability and healthcare utilisation: baseline and instrumented estimates.**

|                                     | (1) Wealth only       | (2) + DCI             | (3) + interaction      | (4) CF-IV              |
|-------------------------------------|-----------------------|-----------------------|------------------------|------------------------|
| Digital Capability Index            |                       | 0.0921***<br>(0.0121) | 0.0911***<br>(0.0121)  | 0.2681***<br>(0.0954)  |
| Wealth rank (centred)               | 1.2723***<br>(0.0356) | 1.1658***<br>(0.0375) | 1.1654***<br>(0.0374)  | 0.9620***<br>(0.1187)  |
| DCI $\times$ wealth rank            |                       |                       | -0.1644***<br>(0.0317) | -0.2553***<br>(0.0466) |
| First-stage residual $\hat{v}$      |                       |                       |                        | -0.1796*<br>(0.0963)   |
| $\hat{v} \times$ wealth rank        |                       |                       |                        | 0.1778***<br>(0.0643)  |
| Implied $\beta_1$ on uncentred rank |                       |                       | 0.1733                 | 0.3957                 |

|                               | (1) Wealth only | (2) + DCI | (3) + interaction | (4) CF-IV |
|-------------------------------|-----------------|-----------|-------------------|-----------|
| Need and demographic controls | Yes             | Yes       | Yes               | Yes       |
| Caste, city-class, state FE   | Yes             | Yes       | Yes               | Yes       |
| R <sup>2</sup>                | 0.228           | 0.231     | 0.232             | 0.233     |
| Observations                  | 14,200          | 14,200    | 14,200            | 14,200    |
| First-stage partial F         |                 |           |                   | 102.9     |
| Over-identification p-value   |                 |           |                   | 0.656     |

**Notes:** Dependent variable is the Healthcare Utilisation Index (0–4). Block-clustered standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ . Column (4) is a control-function estimator using block backhaul connectivity and the leave-one-out neighbour connectivity mean as excluded instruments; the interaction is instrumented through the additional residual interaction term.

Table 5 disaggregates by component. Average marginal effects of digital capability are positive and significant for all four components, but they are largest for preventive screening (0.0283) and qualified-provider contact (0.0158), and the interaction is most negative for preventive screening (-0.0225). This is substantively coherent: preventive care is the margin on which information and scheduling frictions bind hardest, and those frictions are most severe for poorer households.

**Table 5: Component-wise average marginal effects (probit).**

| Component                    | DCI                   | Wealth rank           | DCI × rank             |
|------------------------------|-----------------------|-----------------------|------------------------|
| Qualified-provider contact   | 0.0158***<br>(0.0048) | 0.2400***<br>(0.0153) | -0.0431***<br>(0.0126) |
| Preventive screening         | 0.0283***<br>(0.0056) | 0.3136***<br>(0.0172) | -0.0225<br>(0.0145)    |
| Chronic-condition follow-up  | 0.0209***<br>(0.0052) | 0.2911***<br>(0.0158) | -0.0234<br>(0.0143)    |
| Care not deferred or forgone | 0.0223***<br>(0.0053) | 0.2915***<br>(0.0152) | -0.0331**<br>(0.0135)  |

**Notes:** Average marginal effects from probit models with the full control set; block-clustered standard errors in parentheses.

Secondary outcomes are informative about the mechanism. Digital capability has no detectable effect on hospitalisation (0.0030,  $p = 0.181$ ), which is what one would expect if inpatient admission is driven by acute clinical need rather than by information frictions. It shifts care toward private providers

(0.0213) and modestly reduces log out-of-pocket expenditure conditional on utilisation (-0.0501). Digital access therefore operates on the ambulatory and preventive margin, not the inpatient one.

## 6.2 Returns to digital capability are pro-poor

Figure 5(a) plots the model-implied marginal effect of digital capability by wealth decile. It declines monotonically from 0.383 in the poorest decile to 0.153 in the richest, a fall of 60%. The confidence band is wide, reflecting the variance cost of instrumentation, but the gradient itself is estimated with a  $t$  statistic of 5.5.

Panel (b) reports a non-parametric check. Estimating separate gradient-boosted regressions for above- and below-median digital capability and differencing gives a conditional average treatment effect that is highest in the poorest quintile (0.286) and among the lowest in the richest (0.130). The pattern is not perfectly monotone — the fourth quintile shows an upward departure — which we read as sampling noise in a flexible learner rather than as evidence of a genuine non-monotonicity. The broad direction corroborates the parametric interaction.




**Fig. 5.** Heterogeneous returns to digital capability. Panel (a) plots the model-implied marginal effect from the control-function specification with a 95% band. Panel (b) plots the difference in gradient-boosted predictions between above- and below-median digital capability by wealth quintile; error bars show the dispersion of individual predictions within quintile and are not inferential standard errors.

## 6.3 Inequality in utilisation and in digital access

Table 6 and Figure 4 report concentration indices. Utilisation is pro-rich,  $CI(HUI) = 0.1066$  (SE 0.0021), falling only to 0.0878 after need standardisation, which tells us that the gradient is not an artefact of the rich being sicker or older. Within components, preventive screening is the most pro-rich (0.1580) and qualified-provider contact the least (0.0718).

Inequality in digital access exceeds inequality in utilisation on every comparable measure. Internet use has CI = 0.1135 and digital health service use has CI = 0.3480, three times the index for utilisation itself. The concentration curves in Figure 4(a) make the ordering visible: the curve for digital health service use lies furthest below the diagonal throughout.

**Table 6: Concentration indices by outcome.**

| Variable                     | CI     | Std. error | Erreygers index |
|------------------------------|--------|------------|-----------------|
| Healthcare Utilisation Index | 0.1066 | 0.0021     | 0.2640          |
| HUI, need-standardised       | 0.0878 | 0.0020     | —               |
| Qualified-provider contact   | 0.0718 | 0.0027     | 0.2186          |
| Preventive screening         | 0.1580 | 0.0049     | 0.2912          |
| Chronic-condition follow-up  | 0.1139 | 0.0038     | 0.2712          |
| Care not deferred            | 0.1044 | 0.0033     | 0.2749          |
| Hospitalisation, last year   | 0.1019 | 0.0244     | 0.0161          |
| Private provider use         | 0.0677 | 0.0035     | 0.1767          |
| Internet use, last 3 months  | 0.1135 | 0.0026     | 0.3355          |
| Digital health service use   | 0.3480 | 0.0076     | 0.3480          |

**Notes:** Kakwani convenient-regression standard errors. The Erreygers index rescales the concentration index for bounded variables and is reported where bounds are defined.




**Figure 4.** Wealth-related inequality in digital access and healthcare utilisation. Panel (a) plots concentration curves against the line of equality; greater distance below the diagonal indicates greater pro-rich concentration. Panel (b) reports concentration indices with 95% confidence intervals, with digital indicators in dark blue and utilisation outcomes in green.

#### 6.4 Decomposition

Figure 6(a) reports the standard Wagstaff decomposition. The direct wealth-rank contribution dominates at 0.0783, roughly 73% of the total index of 0.1066. Education contributes 0.0080 and insurance 0.0058. The digital capability contribution computed in the conventional way is only 0.0106, about 10% of the total — a figure that would support the conclusion that digital access is a minor player.

That conclusion is an artefact of the method.

Applying Proposition 1 with the instrumented coefficients and the interaction term carried explicitly raises the digital contribution to 0.0311, a factor of 2.9. Two things drive the gap. The OLS coefficient is attenuated by measurement error in a six-item latent index, and the standard procedure has no place for the interaction, which in our setting contributes a non-trivial offsetting term. This is a general point about decomposition practice, not a feature of our particular data: wherever a regressor is a latent construct measured by a short scale and interacts with the ranking variable, the conventional decomposition will understate its role.

#### 6.5 The access and returns channels

Table 7 and Figure 6(b) report the central result. The access channel is  $A = 0.0456$ : digital capability is pro-rich, with  $\text{Cov}(\theta, r) = 0.1426$ , and it is productive, so its distribution amplifies inequality. The returns channel is  $R = -0.0145$ : because  $\gamma < 0$  and  $\text{Cov}(vr, r) = 0.0702 > 0$ , differential returns work in the opposite direction. But the access channel is roughly 3.1 times larger in magnitude. The net contribution is  $\Delta = 0.0322$  (block-bootstrap 95% CI 0.0109 to 0.0518).

Equalising digital capability at its current mean would therefore reduce the concentration index of utilisation from 0.1066 to 0.0744. Digital stratification accounts for  $S = 30.2\%$  of pro-rich inequality in utilisation (95% CI 10.2–48.8%), a Digital Inequality Multiplier of  $M = 1.43$  (95% CI 1.11–1.95). The corresponding OLS-based figures are markedly smaller ( $S = 10.8\%$ ,  $M = 1.12$ ), and we regard them as a lower bound for the reasons given above.

Proposition 2 supplies the policy-relevant threshold. Given the estimated  $\beta_1$  and the observed distribution of digital capability, returns would have to satisfy  $\gamma < \gamma^* = -0.803$  for digital access to be inequality-neutral. The estimated  $\gamma$  is  $-0.255$ . The equity gap is  $G = 0.548$ : returns are pro-poor, but they would need to be about 3.1 times more steeply pro-poor to offset the unequal distribution of access. Figure 7(b) plots the affine relationship and marks both values.

**Table 7: Access–returns decomposition and the Digital Inequality Multiplier.**

| Quantity                            | Symbol                  | Estimate | 95% CI |
|-------------------------------------|-------------------------|----------|--------|
| Productivity of digital capability  | $\beta_1$               | 0.3957   | —      |
| Interaction with wealth rank        | $\gamma$                | -0.2553  | —      |
| Concentration of digital capability | $\text{Cov}(\theta, r)$ | 0.1426   | —      |

| Quantity                             | Symbol                  | Estimate | 95% CI             |
|--------------------------------------|-------------------------|----------|--------------------|
| Second-order covariance              | $\text{Cov}(v_r, r)$    | 0.0702   | —                  |
| Access channel                       | A                       | 0.0456   | [0.0233, 0.0677]   |
| Returns channel                      | R                       | -0.0145  | [-0.0193, -0.0097] |
| Net contribution                     | $\Delta = A + R$        | 0.0322   | [0.0109, 0.0518]   |
| Observed concentration index         | CI(Y)                   | 0.1066   | —                  |
| Counterfactual concentration index   | CI(Y <sup>cf</sup> )    | 0.0744   | —                  |
| Share attributable to digital access | S                       | 30.2%    | [10.2%, 48.8%]     |
| Digital Inequality Multiplier        | M                       | 1.433    | [1.114, 1.955]     |
| Neutrality threshold                 | $\gamma^*$              | -0.803   | [-1.196, -0.409]   |
| Equity gap                           | $G = \gamma - \gamma^*$ | 0.548    | —                  |

**Notes:** Estimates use the control-function coefficients. Confidence intervals from 500 block bootstrap replications resampling primary sampling units, re-estimating the first and second stages and recomputing wealth ranks within each replicate.




**Figure 6.** Decomposing wealth-related inequality in healthcare utilisation. Panel (a) reports the standard Wagstaff decomposition; the hatched bar shows the digital contribution recomputed with instrumented coefficients and the interaction term carried explicitly. Panel (b) is a waterfall representation of Proposition 1: removing the access channel lowers the concentration index, removing the returns channel raises it, and the counterfactual index is the net.

6.6 Counterfactual policy simulations

Figure 7(a) places six scenarios in mean-utilisation and inequality space. The observed point sits at mean HUI 2.475 with CI 0.1065. A universal one-standard-deviation increase in digital capability delivers the largest gain in mean utilisation (to 2.743) but only a partial reduction in inequality (to 0.0805). A one-standard-deviation increase confined to the bottom two quintiles delivers about 53% of the mean gain while producing a strictly lower concentration index (0.0680) than the universal policy — achieved with

two-fifths of the reach.

The mirror-image scenario is instructive. Extending the same increase to the top three quintiles raises mean utilisation by a comparable amount but pushes the concentration index up to 0.1179, above the observed level. An untargeted programme that in practice reaches the connected first — which is what universal digital-health rollouts tend to do — is not merely less equalising than a targeted one. It actively widens the gradient while improving the average.

**Table 8. Counterfactual simulations of digital capability policies.**

| Scenario                                       | Mean HUI | $\Delta$ mean | CI     | $\Delta$ CI |
|------------------------------------------------|----------|---------------|--------|-------------|
| Observed                                       | 2.475    | —             | 0.1065 | —           |
| Full equalisation of $\theta$ at constant mean | 2.512    | +0.037        | 0.0744 | -0.0322     |
| Bottom 40% raised to the median                | 2.585    | +0.110        | 0.0736 | -0.0330     |
| Universal +1 SD                                | 2.743    | +0.267        | 0.0805 | -0.0261     |
| Bottom 40% +1 SD (targeted)                    | 2.616    | +0.141        | 0.0680 | -0.0386     |
| Top 60% +1 SD                                  | 2.602    | +0.127        | 0.1179 | +0.0113     |

**Notes:** Predicted from the control-function specification, holding all other covariates at observed values and truncating predictions to the admissible range [0,4].




**Figure 7.** Policy simulations and the neutrality threshold. Panel (a) locates each scenario in mean-utilisation and inequality space; the targeted scenario (5) dominates the universal scenario (4) on the equity dimension at a fraction of the reach. Panel (b) plots the net contribution  $\Delta$  as an affine function of the interaction coefficient  $\gamma$ , marking the estimated value, the neutrality threshold  $\gamma^*$ , and the equity gap between them.

6.7 Gender

Figure 8(a) shows that the male–female gap is narrow for device ownership and much wider for the skill-based competences. Phone ownership differs by 6 percentage points, whereas the ability to

send an attachment differs by 26 points and online transactions by 25 points. The gap narrows somewhat for digital health service use (19 points), but only because prevalence is low for both sexes. Since it is the higher-order competences that carry the utilisation effect, a device-ownership metric would substantially understate the gendered component of digital exclusion. This mirrors NFHS-5, where the urban gender gap in ever having used the internet is roughly twenty percentage points.

## 7. VALIDATION, ROBUSTNESS AND LIMITATIONS

### 7.1 Monte Carlo validation of the estimator

Because the data-generating process is known, we can ask directly whether the estimator of Proposition 1 recovers the true  $\Delta$ . We conduct 60 independent replications of a reduced version of the design ( $n = 5,000$ , 220 blocks). In each replication we compute the true  $\Delta$  by re-drawing outcomes under the counterfactual using common random numbers, and compare it with two estimators: the control-function estimator used throughout, and a naive OLS version. The results, in Table 9 and Figure 8(b), are informative in both directions. The true mean  $\Delta$  is 0.0214. The naive OLS estimator averages 0.0104, a bias of -0.0110 — it recovers less than half of the true contribution, confirming the attenuation argument of Section 6.4. The control-function estimator averages 0.0266, a bias of 0.0052, roughly 24% upward. It is much closer to the truth in expectation but noticeably noisier (RMSE 0.0227 against 0.0117).

We report this candidly rather than presenting the instrumented estimate as unbiased. The honest summary is that OLS reliably understates the digital contribution by roughly half, the control-function estimator overstates it by roughly a quarter, and the truth lies between them but closer to the instrumented figure. Our headline share of 30% should accordingly be read as an upper-middle estimate within a plausible range of roughly 15% to 30%.

**Table 9: Monte Carlo validation of the decomposition estimator (60 replications).**

| Estimator                  | Mean $\Delta$ | Bias    | RMSE   | Relative bias |
|----------------------------|---------------|---------|--------|---------------|
| True $\Delta$ (known DGP)  | 0.0214        | —       | —      | —             |
| Control-function estimator | 0.0266        | 0.0052  | 0.0227 | +24.1%        |
| Naive OLS estimator        | 0.0104        | -0.0110 | 0.0117 | -51.4%        |

**Notes:** Each replication regenerates the population, refits the item response model, re-estimates first and second stages,

and recomputes the decomposition. The true  $\Delta$  is obtained by simulating the counterfactual with common random numbers.




**Figure 8.** Gender and validation. Panel (a) shows that the male–female gap is narrow for device ownership and much wider for skill-based competences; figures above bars are percentage-point gaps. Panel (b) compares the true  $\Delta$  with the control-function and naive OLS estimators across Monte Carlo replications.

### 7.2 Robustness

The substantive conclusions survive a range of alternative choices. Using the Erreygers-corrected index in place of the standard concentration index changes the levels but not the ordering or the sign of any channel. Need standardisation reduces CI(HUI) from 0.1066 to 0.0878 without altering the decomposition qualitatively. Replacing the item response score with a simple unweighted sum of the six indicators attenuates  $\beta_1$  as expected but leaves the sign and approximate magnitude of  $\gamma$  intact. Dropping either instrument individually leaves the first-stage F above conventional thresholds and moves the point estimates within their confidence intervals.

### 7.3 Limitations

**The data are synthetic.** This is the binding limitation and it constrains the interpretation of every number in the paper. The dataset reproduces published aggregate margins from NFHS-5, NSS 75th Round and NSS 79th Round CAMS, but it is generated from a model rather than sampled from the Indian population. It cannot support claims about the actual magnitude of digital-related inequality in India. What it can and does support is a methodological claim: that the access–returns decomposition is well defined, that it is estimable, that its behaviour under known truth can be characterised, and that the conventional decomposition understates the digital contribution in an identifiable way. Applying the method to NSS 75th and 79th Round unit records, available from the MoSPI microdata portal, and to NFHS-5 records from the DHS Program, is the natural next step and

the one we would prioritise.

**The exclusion restriction is imposed, not tested.**

In the calibrated design, block connectivity affects utilisation only through digital capability by construction. In real data this would need defending. Telecom rollout may correlate with municipal investment that also improves health infrastructure, and the leave-one-out instrument is subject to the reflection problem. Both concerns would require institutional argument and additional falsification tests.

**The linear-interaction functional form is restrictive.**

Proposition 1 is exact within model (1) but the true relationship between digital capability, wealth and utilisation may be non-linear in ways a single interaction cannot capture. The gradient-boosted results in Figure 5(b) are broadly consistent with the parametric specification but not perfectly monotone. A more general treatment would state the decomposition for an arbitrary conditional mean function, at the cost of the closed-form neutrality threshold.

**Utilisation is not health.** The HUI measures contact with the health system, weighted equally across four components. Equal weighting is a normative choice, and more utilisation is not unambiguously better — some of the private-provider shift we detect may reflect supplier-induced demand rather than met need. Linking digital access to outcomes rather than to contact is an important extension.

**Cross-sectional design.** Digital capability and healthcare utilisation are measured contemporaneously, so reverse causation — illness episodes prompting digital adoption — cannot be excluded beyond what the instruments handle.

## 8. DISCUSSION AND POLICY IMPLICATIONS

The central finding is a conflict between two facts that are individually well documented and jointly under-appreciated. Digital access does more good for poorer households, and poorer households have much less of it. The decomposition of Section 3 shows that these are not competing intuitions to be adjudicated rhetorically but additive terms whose relative magnitudes can be measured. In our calibrated setting the second dominates the first by a factor of roughly 3.1.

Three implications follow. First, the case for targeting

is stronger than the case for universality, and the argument is not the usual one about cost-effectiveness under a budget constraint. Because returns are pro-poor, targeting the bottom of the distribution is more efficient per unit delivered, not merely more equitable. Our simulations show a targeted intervention reaching two-fifths of the population achieving a lower concentration index than a universal one, while capturing a substantial share of the mean gain. Universal digital-health programmes that in practice reach the already-connected first can widen the gradient even as they raise the average — an outcome that aggregate utilisation statistics will not reveal.

Second, digital-health programme monitoring should track the concentration index of platform use, not just its volume. A platform reporting hundreds of millions of consultations may be reducing or increasing inequality depending entirely on who is consulting. The eSanjeevani architecture is instructive here: the provider-assisted health and wellness centre model, in which a frontline worker mediates the consultation, decouples the benefit from the patient's own digital capability, whereas the direct-to-patient model does not. Our framework implies these two modalities have opposite distributional signatures, and that the mediated model is the equalising one.

Third, investment should target the competences that carry the effect rather than device penetration. Our data show phone ownership near saturation across the wealth distribution while the ability to complete a digitally mediated health interaction remains sharply stratified. Handset subsidies address a margin that is nearly closed; digital-literacy provision, vernacular interfaces, and assisted-access points address the margin that is not. The gender pattern reinforces this: the male-female gap is small in device ownership and large in the higher-order competences.

Finally, a methodological implication for the health-equity literature. Analysts routinely decompose concentration indices using OLS coefficients and omit interactions. Where the regressor of interest is a latent construct measured by a handful of binary indicators — which describes not only digital access but also health literacy, social capital, and empowerment indices — this practice will systematically understate its contribution. Our Monte Carlo puts the understatement at roughly half. That is large enough to change conclusions.

## 9. CONCLUSION

We have asked whether digital access equalises healthcare utilisation in urban India and have supplied the formal apparatus needed to give the question a determinate answer. The contribution of digital access to wealth-related inequality in utilisation decomposes exactly into an access channel and a returns channel, and there exists a unique threshold value of the interaction coefficient at which the two cancel. Applying this to a dataset calibrated to national survey margins, we find that returns to digital capability are markedly pro-poor — falling by 60% from the bottom decile to the top — yet the pro-rich distribution of that capability dominates, so that digital stratification accounts for roughly 30% of measured inequality in utilisation. Returns would need to be about 3.1 times more steeply pro-poor for digital access to be distributionally neutral.

The result is not an argument against digital health. It is an argument that the distributional consequences of digital health depend on a quantity — the covariance of capability with socioeconomic rank — that programme design can influence and that programme evaluation currently ignores. A digital health system can be made an equaliser. It will not become one by default.

Two extensions seem most valuable. The first is application to survey unit records, which would convert the calibrated magnitudes here into population estimates. The second is a semiparametric version of Proposition 1 that dispenses with the linear interaction while retaining an interpretable neutrality condition.

## REFERENCES

- Andersen, R. M. (1995). Revisiting the behavioral model and access to medical care: does it matter? *Journal of Health and Social Behavior*, 36(1), 1–10.
- Athey, S., & Imbens, G. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360.
- Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: application of an EM algorithm. *Psychometrika*, 46(4), 443–459.
- DiMaggio, P., & Hargittai, E. (2001). From the “digital divide” to “digital inequality”: studying Internet use as penetration increases. Working Paper 15, Center for Arts and Cultural Policy Studies, Princeton University.
- Embretson, S. E., & Reise, S. P. (2000). *Item Response Theory for Psychologists*. Mahwah, NJ: Lawrence Erlbaum.
- Erreygers, G. (2009). Correcting the concentration index. *Journal of Health Economics*, 28(2), 504–515.
- Filmer, D., & Pritchett, L. H. (2001). Estimating wealth effects without expenditure data—or tears: an application to educational enrollments in states of India. *Demography*, 38(1), 115–132.
- Hargittai, E. (2002). Second-level digital divide: differences in people's online skills. *First Monday*, 7(4).
- International Institute for Population Sciences (IIPS) & ICF (2021). *National Family Health Survey (NFHS-5), 2019–21: India Report*. Mumbai: IIPS.
- Kakwani, N., Wagstaff, A., & van Doorslaer, E. (1997). Socioeconomic inequalities in health: measurement, computation, and statistical inference. *Journal of Econometrics*, 77(1), 87–103.
- Ministry of Health and Family Welfare, Government of India (2020). *Telemedicine Practice Guidelines*. New Delhi.
- Ministry of Statistics and Programme Implementation (2019). *Health in India. NSS 75th Round, July 2017–June 2018, Report No. 586*. New Delhi: National Statistical Office.
- Ministry of Statistics and Programme Implementation (2024). *Comprehensive Annual Modular Survey, NSS 79th Round, July 2022–June 2023*. New Delhi: National Statistical Office.
- Norris, P. (2001). *Digital Divide: Civic Engagement, Information Poverty, and the Internet Worldwide*. Cambridge: Cambridge University Press.
- O'Donnell, O., van Doorslaer, E., Wagstaff, A., & Lindelow, M. (2008). *Analyzing Health Equity Using Household Survey Data*. Washington, DC: World Bank.
- Rivers, D., & Vuong, Q. H. (1988). Limited information estimators and exogeneity tests for simultaneous probit models. *Journal of Econometrics*, 39(3), 347–366.
- Robinson, L., Cotten, S. R., Ono, H., Quan-Haase, A., Mesch, G., Chen, W., Schulz, J., Hale, T. M., & Stern, M. J. (2015). Digital inequalities and why they matter. *Information, Communication & Society*, 18(5), 569–582.
- Selvaraj, S., & Karan, A. K. (2009). Deepening health insecurity in India: evidence from national sample surveys since 1980s. *Economic and Political Weekly*, 44(40), 55–60.
- Sood, S., Lal, K., Bhatia, M., Kapoor, G., Singh, S., Kaushish, R. K., et al. (2025). Adoption and utilization of India's eSanjeevani national telemedicine service. *Oxford Open Digital Health*, 3, oqaf025.
- van Dijk, J. A. G. M. (2005). *The Deepening Divide: Inequality in the Information Society*. Thousand Oaks, CA: Sage.
- Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523), 1228–1242.
- Wagstaff, A. (2005). The bounds of the concentration index when the variable of interest is binary, with an application to immunization inequality. *Health Economics*, 14(4), 429–432.
- Wagstaff, A., van Doorslaer, E., & Watanabe, N. (2003). On decomposing the causes of health sector inequalities with an application to malnutrition inequalities in Vietnam. *Journal of Econometrics*, 112(1), 207–223.
- Wooldridge, J. M. (2015). Control function methods in applied econometrics. *Journal of Human Resources*, 50(2),

## APPENDIX A. SUPPLEMENTARY DERIVATIONS

### A.1 The second-order covariance under a linear capability gradient

The returns channel depends on  $\text{Cov}(vr, r)$ , which is not a familiar object. It is useful to evaluate it under a simple approximation. Suppose the deviation of digital capability from its mean is linear in rank,  $v = c(r - 1/2) + e$  with  $e$  independent of  $r$  and mean zero, where  $c > 0$  measures how steeply capability rises with wealth. Then, using  $E[r^2] = 1/3$  and  $E[r^3] = 1/4$ ,

$$E[vr^2] = c(E[r^3] - E[r^2]/2) = c(1/4 - 1/6) = c/12,$$

$$E[vr] = c(E[r^2] - E[r]/2) = c(1/3 - 1/4) = c/12,$$

so that  $\text{Cov}(vr, r) = E[vr^2] - E[vr]E[r] = c/12 - c/24 = c/24$ . Two consequences follow. The second-order covariance is positive whenever capability is pro-rich, confirming the premise of Corollary 1. And since  $\text{Cov}(\theta, r) = c\text{Var}(r) = c/12$  under the same approximation, the neutrality threshold simplifies to

$$\gamma^* \approx -\beta_1(c/12)/(c/24) = -2\beta_1,$$

independent of  $c$ . This is a striking simplification: to first order, digital access is inequality-neutral precisely when the marginal return falls to zero at the top of the distribution, since  $\beta_1^u + \gamma = \beta_1^u - 2\beta_1^u < 0$ . With  $\beta_1 = 0.396$  the approximation gives  $\gamma^* \approx -0.791$ , close to the exact value of  $-0.803$ ; the discrepancy reflects departures from a strictly linear capability gradient.

### A.2 Centred and uncentred parameterisations

Estimation uses the centred rank  $r^c = r - 1/2$  so that the coefficient on  $\theta$  is the marginal effect at median wealth. Since  $\beta_1^c\theta + \gamma\theta r^c = (\beta_1^c - \gamma/2)\theta + \gamma\theta r$ , the uncentred productivity parameter needed for Proposition 1 is  $\beta_1^u = \beta_1^c - \gamma/2$  while  $\gamma$  is invariant. Failing to apply this transformation causes the returns channel to vanish identically whenever the capability gradient is linear in rank, because  $\text{Cov}(\theta r^c, r) - \theta\text{Var}(r) = \text{Cov}(vr, r) - \text{Cov}(v, r)/2 = c/24 - c/24 = 0$ . This is an easy error to make and it silently removes the entire equalising channel from the decomposition.

### A.3 Delta-method variance of the net contribution

Treating the covariances and  $\mu$  as fixed at their sample values, the absolute net contribution  $A_a + R_a$  is linear in  $(\beta_1^c, \gamma)$  with gradient  $J = 2[\text{Cov}(\theta, r), \text{Cov}(vr, r) - \text{Cov}(\theta, r)/2]/\mu$ , giving  $\text{Var} = J'VJ$  where  $V$  is the block-

clustered covariance of the two structural coefficients. This yields a standard error of 0.0110 on the absolute scale. We prefer the block bootstrap reported in the main text because it also propagates uncertainty in the estimated ranks and in the item response scores.

## APPENDIX B. WEALTH INDEX CONSTRUCTION

The wealth index is the first principal component of ten binary asset and housing indicators, standardised before extraction. The component explains 38.2% of total variance. Table B1 reports the loadings.

**Table B1. First principal component loadings for the asset-based wealth index.**

| Indicator                 | Loading | Indicator       | Loading |
|---------------------------|---------|-----------------|---------|
| Refrigerator              | 0.374   | Car             | 0.314   |
| Washing machine           | 0.371   | Air conditioner | 0.306   |
| Two-wheeler               | 0.363   | Television      | 0.265   |
| Piped water supply        | 0.334   | LPG connection  | 0.247   |
| Pucca (permanent) housing | 0.328   | Bank account    | 0.218   |

**Notes:** Loadings on the first principal component, oriented so that higher values indicate greater wealth. Variance explained: 38.2%.

## APPENDIX C. REPRODUCIBILITY

The analysis comprises four stages: dataset construction; estimation (item response model, principal components, regressions, control-function instrumental variables, concentration indices and the decomposition); Monte Carlo validation; and figure generation. All estimation was carried out in Python 3.12 using NumPy 2.4, pandas 3.0, SciPy 1.17, statsmodels 0.15 and scikit-learn 1.8. The item response model was fitted by custom maximum a posteriori code using Gauss–Hermite quadrature rather than an off-the-shelf package, so that the prior specification and the expected a posteriori scoring are fully transparent. Random seeds are fixed throughout; the block bootstrap uses 500 replications and the Monte Carlo study 60 replications.