Bounding the Uncounted: Partial Identification of Non-Pecuniary Welfare from Innovation, Net of Quality Adjustment

Do the non-pecuniary benefits of innovation — time saved, reliability, reduced cognitive effort, option value — imply that conventional estimates of the social return to R&D are downward biased? We argue the popular affirmative is not identified, and we replace it with a smaller claim that survives its strongest objection. Two facts discipline the exercise. First, the "social rate of return to R&D" in the literature it would amend (Jones and Williams, 1998, QJE; Bloom, Schankerman and Van Reenen, 2013, Econometrica) is a production-side object — the elasticity of output or TFP with respect to R&D capital, with the social–private wedge arising from knowledge spillovers net of business-stealing. Consumer time saved is consumer surplus, a different estimand; attaching it to the production-side return is a category error. Second, official statistics already absorb part of what looks "non-pecuniary": hedonic and quality-adjusted price indices value characteristics such as speed and reliability, and Bils and Klenow (2001, AER) find that official prices do not fully net out quality change (annual quality growth near 3.7% for their goods is only partly removed). The captured share, $\varphi \in (0,1)$, is generally unknown — so the sign of any bias in measured returns is not known a priori, and crediting the same attribute again risks double counting and upward, not downward, bias.

We therefore study the uncounted component directly. Write the social return as $\rho = \rho_{measured} + \Delta$, with $\Delta = (1-\varphi)\cdot v \cdot \Delta T \cdot N / R\&D$ the per-R&D value of an attribute improvement ΔT delivered to N users at shadow value v, and φ the share already embedded in quality-adjusted output. Monotone treatment response (Manski, 1997, Econometrica) — adoption, by revealed preference, does not lower a user's welfare — signs the welfare term, $\Delta \geq 0$, but its sharp lower bound is generically zero: monotonicity yields a direction, not a magnitude. An informative, strictly positive floor requires three further ingredients, each a maintained assumption rather than a consequence of the data: a defended value-of-time floor $v_{LB} > 0$; a documented hedonic gap $\bar{\varphi} < 1$ for the specific attribute; and clean attribution of R&D to that attribute. Then $\Delta \geq (1-\bar{\varphi})\cdot v_{LB} \cdot \Delta T \cdot N / R\&D > 0$. The upper bound is $+\infty$ unless one also imposes a finite ceiling v_{UB} and $\varphi_{LB} > 0$; "bounded magnitude" is otherwise an overstatement.

Stated this way, the contribution is as much a map of when the non-pecuniary term can be signed as a claim that it is — and the conditions rarely hold jointly. The value-of-time floor is fragile: the wage bounds the value of saved time only if the freed time is reallocated to uses worth at least the wage, yet time-diary evidence (Aguiar and Hurst, 2007, QJE) shows freed time often flows to low-valued leisure, and the shadow price is a Lagrange multiplier on a binding time constraint (Becker, 1965; DeSerpa, 1971), not a "no-arbitrage" price recoverable from scanner data, which record prices and quantities but no time attribute. The hedonic gap must be shown, not assumed: reliability and speed are frequently inside the price vector already, so the cleanest qualifying attribute is consumer time, which is not transacted and so sits structurally outside it. Attribution is the binding empirical constraint: firm- and industry-level R&D (Compustat, BERD) cannot be apportioned to a single attribute, so the denominator is defined only for a single-purpose program. And even where Δ is signed and bounded, it does not sign the bias in measured social returns: φ is unknown, and social offsets — induced demand and congestion (Duranton and Turner, 2011, AER, estimate a road-travel elasticity near one) and attention as a scarce input — can drive the social non-pecuniary term to zero or below even when each adopter's private term is non-negative.

We illustrate in the one setting where attribution is clean: GPS, a single-purpose public technology whose program R&D is on record and whose civilian time-saving margin is plausibly outside official quality adjustment. (The documented order of magnitude of GPS economic benefits is large — NIST/RTI, 2019 — but we use it only to anchor a transparent lower bound on the uncounted consumer-welfare component, not as a correction to the social R&D return.) The exercise yields a floor that is positive under its maintained assumptions and explicitly silent on the direction of bias in measured returns.

The result is conditional and partly negative by design. Partial identification converts a sign-unknown nuisance into a signed, bounded consumer-welfare term only when an attribute is provably outside quality adjustment, carries a defensible value-of-time floor, and admits clean R&D attribution — a knife-edge met by time-saving from single-purpose technologies and few other cases. The honest headline is not that R&D returns are understated, but that the missing non-pecuniary welfare can be bounded away from zero exactly where these three conditions bind, and not otherwise.

Acknowledgements: This research was conducted independently with no external financial support. The author declares no conflicts of interest.