The Demand Side

Identification error in child-safety markets, and the object that was never offered

Working paper · Ninth revision · August 2026

An inference that does not follow

Child-safe infrastructure is generally held to have failed on the supply side. A prior failure is one of identification. Demand has been inferred from markets in safety instruments — controls, filters, monitoring, settings — and weak uptake of those instruments has been read as weak demand for child safety.

A market for instruments measures three things at once: willingness to be protected, willingness to perform the implementation labour, and capacity to perform it. It reports the product as though it were the first alone.

This paper does not claim latent demand is large. It claims the evidence usually offered against it is non-responsive to the question, that the relevant variable has never been isolated because the relevant object has never existed, and that the object can be specified precisely enough to be built and tested.

What non-adoption does not reveal

The converse matters as much. A household declining a technical instrument has not thereby revealed indifference to the protective end. It may have rejected the instrument, its implementation burden, its restrictions, its surveillance properties or its price; or judged the risk not to warrant it; or preferred rules and conversation, which comparative studies often find more prevalent than technical-control use. Observed non-adoption alone does not distinguish among those explanations, and that is the whole of the identification problem.

Note what the argument does not require. It does not require showing that implementation labour causes non-adoption in any particular household. It requires only that adoption be jointly determined by protective preference and other variables — which the empirical literature establishes directly.

The comparison that makes it plain

Nobody measures demand for safe drinking water by counting sales of domestic filtration units, or demand for food safety by counting home pathogen test kits. In each case the small domestic instrument market is read as evidence that the environment is governed upstream, not as evidence of an indifferent population. Where such markets are large — bottled water on an unreliable supply — the size is read correctly, as institutional failure rather than consumer enthusiasm.

The child-safety instrument market has been read the other way round.

What the object would have to be

Six properties follow, and they are not six predictions about what households value. They are controls on what an election is allowed to mean: selectable in a single decision, portable across ordinary devices and networks, legible in its published conditions, revocable without a designed penalty, calibrated on a schedule the household controls, and not dependent on surveillance of the child.

Each removes one contaminant from the measurement. The experiment then asks a clean question — given an object in which those confounds have been reduced, do households elect it and keep it?

The properties are conditions, not an architecture. More than one construction could satisfy them, and nothing in the argument depends on which. A reader wishing to defeat the paper must attack the conditions.

Two hypotheses, kept apart

The identification hypothesis asks whether transferring safety-management labour upstream materially changes household election and retention. That is the thesis. The supply hypothesis asks whether an observable enrolled audience induces commercial supply. That is a prediction about a two-sided market, inherited from an existing literature.

The supply hypothesis can fail while the thesis holds. That outcome would establish that households want the object and that markets will not spontaneously produce one — a finding about how such an environment would have to be financed, not a refutation. The paper is organised so that result can be reached and reported as such.

What is not known, and what would settle it

Three problems are most likely to be decisive: whether interpersonal communication can cross the boundary without either admitting unapproved contact or restoring configuration labour; who pays for an environment that refuses to make the child an advertising product; and whether a household election can be expressed without producing a persistent child identifier.

The paper does not claim any coalition exists, does not claim regulatory endorsement, and proposes compulsion of no party. What it offers is that the not-knowing is a solved problem rather than a permanent condition: the object can be specified, an election is a countable act, and a controlled trial against a configured tool at comparable coverage and varied price can answer the question in either direction.

The full paper sets out the empirical programme, the conditions under which the argument should be rejected, and the trial design in detail. The author holds a commercial interest in technology related to the architecture used in the paper as an illustration, through an entity separate from the foundation; the argument does not depend on that architecture, and the paper says why.

To request the full paper, write to shknudson@ifcsis.org.