Comparing the Standard of Living: A Nonparametric Approach to the Cost-of-Living Index

Davide M. Difino
Goethe University Frankfurt & GSEFM

Measuring the material standard of living in a comparable way remains a central challenge in economics. It is well known that conventional price adjustments impose restrictive assumptions on preferences, ignoring changes over time, income, age, and other socio-demographic characteristics that shape them. This project develops a nonparametric algorithm to approximate the ideal Cost-of-Living Index, expressing compensated expenditure shares as functions of observed shares and (compensated) price elasticities, which are estimable from panel data. The method accommodates arbitrary non-homothetic and state-variable-driven preferences. Applying this approach to multi-country household data enables tracing the evolution of the standard of living across groups and over time.

The measurement of the material standard of living in a way that permits meaningful comparison is among the oldest and most persistent challenges in economics. It is a prerequisite for virtually every question we ask – positive or normative – about growth, policy, and development. Yet despite its centrality, we lack an empirical measure of the standard of living that is theoretically coherent, empirically sound, and practically implementable. The most common practical approach has long been money-metric: we compare expenditures expressed in a common monetary unit. The appeal is immediate. Such measures are observable, comparable, and grounded in behavior. But money is not a neutral metric of the standard of living. What a given amount of money can buy changes with prices, and what that consumption represents in terms of the standard of living depends on preferences.

The issue that the purchasing power of money and the material standard of living it represents are contingent upon prices and preferences has been addressed in practice through deflation using price indexes. These indexes, such as the Laspeyres, Paasche, Fisher, or Törnqvist, approximate the continuous Divisia index and are used by statistical agencies to track the evolution of consumer prices over time. Yet every such index rests on strong and often implicit assumptions about preferences. They presume that consumers share identical, time-invariant preferences unaffected by income, age, household composition, or other socio-demographic characteristics (e.g. Baqaee and Burstein, 2023). In effect, they assume that an 80-year-old pensioner and a 20-year-old student would display identical expenditure patterns if faced with the same prices and income – a convenient but empirically untenable fiction.

If we wish to make meaningful comparisons of the standard of living when income, tastes, or other characteristics change, we must appeal to a deeper construct: the Cost-of-Living Index. The Cost-of-Living Index asks how much expenditure would be required to maintain a constant standard of living across different states of the world. It captures the idea of equivalence in the standard of living rather than equality in nominal expenditure. The conceptual difficulty, however, is immediate. A true Cost-of-Living Index requires knowledge of an inherently counterfactual quantity – the expenditure necessary to achieve the same standard of living under alternative price vectors. To compute this, one must know the underlying utility function, which is not directly observable. Under restrictive assumptions – such as homothetic CES preferences identical across agents – the Cost-of-Living Index simplifies and coincides with the familiar consumer price index.

This observation motivates the central question of this project: how can we measure a Cost-of-Living Index without making functional-form assumptions on preferences? The project proposes a new algorithm that makes an approximation possible. The starting point is the recognition that the theoretically ideal Cost-of-Living Index – formally, the Konüs price index – depends on compensated expenditure shares, which are not directly observable. First, I show that these compensated shares can be expressed as functions of observed expenditure shares and (compensated) price elasticities for any arbitrary preference structure. Second, I show that price elasticities can be directly estimated from panel data, allowing us to approximate the Cost-of-Living Index empirically without assuming a specific functional form for preferences. The resulting measure remains grounded in observed behavior while retaining theoretical validity under broad conditions.

This approach has several advantages relative to recent contributions in the literature. For instance, Redding and Weinstein (2020) compute changes in the standard of living under CES preferences with idiosyncratic taste shocks, allowing preferences to evolve over time but only within a fixed functional structure. My framework dispenses with that structure entirely, also accommodating preferences in which consumption patterns vary systematically with income (i.e. non-homothetic preferences). Similarly, while Jaravel and Lashkari (2024), and Baqaee et al. (2024) propose non-parametric corrections to conventional price indexes to capture non-homotheticity, my method extends this flexibility to preferences that vary with any observable socio-demographic characteristics such as age, household composition, education, or location. These extensions are empirically motivated: prior research demonstrates that such characteristics exert strong and systematic effects on preferences. Moreover, the estimation of price elasticities within relatively homogeneous subpopulations allows for direct comparisons of the standard of living across demographic groups, rather than only within them.

Having established a framework capable of producing group-specific Cost-of-Living Indexes, the next step follows naturally. If we can compute a theoretically valid Cost-of-Living Index for each socio-demographic group, we can trace the evolution of the standard of living across these groups over time. This leads to the second guiding question of the project: How has the material standard of living evolved across socio-demographic groups? To answer it, I will apply the algorithm to household panel data from multiple countries. The resulting group-specific Cost-of-Living Indexs will permit comparisons of the standard of living across income levels, age cohorts, and household types, revealing how differences in preferences and price responses shape its trajectory.

In summary, this project develops a general and empirically implementable framework for the measurement of the standard of living. By linking the theoretical definition of the Cost-of-Living Index to observable data without imposing functional-form assumptions, it unifies the theoretical rigor of economic measurement with empirical tractability. The broader contribution is twofold: it refines our capacity to measure changes in the standard of living over time, and it provides a foundation for analyzing differences in the standard of living across the heterogeneous populations that constitute modern economies.

References

David R Baqaee and Ariel Burstein. Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics*, 138(2):769–834, 2023.

David R Baqaee, Ariel T Burstein, and Yasutaka Koike-Mori. Measuring welfare by matching households across time. *The Quarterly Journal of Economics*, 139(1):533–573, 2024.

Xavier Jaravel and Danial Lashkari. Measuring growth in consumer welfare with income-dependent preferences: nonparametric methods and estimates for the united states. *The Quarterly Journal of Economics*, 139(1):477–532, 2024.

Stephen J Redding and David E Weinstein. Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics*, 135(1):503–560, 2020.