AMERICAN UNIVERSITY OF BEIRUT

MEASURING INEQUALITY AND POVERTY WHEN INCOME DATA ARE MISSING: CAN HOUSEHOLD ASSETS SERVE AS PRACTICAL PROXIES?

by
DANA ISSAM DGHAIDY

A thesis
submitted in partial fulfillment of the requirements
for the degree of Master of Arts
to the Department of Economics
of the Faculty of Arts and Sciences
at the American University of Beirut

Beirut, Lebanon
May 2026

AMERICAN UNIVERSITY OF BEIRUT

MEASURING INEQUALITY AND POVERTY WHEN INCOME DATA ARE MISSING: CAN HOUSEHOLD ASSETS SERVE AS PRACTICAL PROXIES?

by
DANA ISSAM DGHAIDY

Approved by:

Dr. Abdallah Zalghout, Assistant Professor Department of Economics	Advisor
Dr. Ali Abboud, Assistant Professor Department of Economics	Member of Committee
Dr. Zeeshan Samad, Assistant Professor Department of Economics	Member of Committee

Date of thesis defense: May 7, 2026

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my thesis advisor, Dr. Abdallah Zalghout, for his guidance, support, and valuable feedback throughout the development of this thesis.

I am eternally grateful to the professors, instructors, and colleagues I had the privilege to cross paths with at the Department of Economics. One may meet brilliant people every day, but it is rare to find people who combine intellect with integrity, respect, and authenticity. I am thankful for the lessons, conversations, and encouragement that shaped both my academic and personal growth.

To my parents: I am nothing without you. I owe it all to you. Your love, sacrifices, and constant belief in me have carried me through every stage of this journey.

To AUB: I hope I gave you even half of what you gave me.

ABSTRACT

OF THE THESIS OF

Dana Issam Dghaidy for Master of Arts
Major: Economics

Title: Measuring Inequality And Poverty When Income Data Are Missing: Can Household Assets Serve As Practical Proxies?

This thesis investigates whether administrative records of housing and vehicles can serve as reliable proxies for measuring economic inequality and poverty in data-scarce environments. Using data from the United States (PSID/SCF) and Indonesia (IFLS), the study evaluates these proxies against a comprehensive suite of indices, including the Gini, Atkinson, and FGT measures. The results demonstrate that asset-based distributions, particularly vehicle ownership, carry a genuine distributional signal that closely tracks income inequality and reacts consistently to macroeconomic shifts like changes in unionization and unemployment. While the findings show that these proxies effectively capture broad social stratification, they also highlight a structural limitation in capturing the extreme upper tail where wealth shifts from physical assets to financial instruments. Ultimately, the study concludes that administrative asset registries can offer a robust, low-cost, and independent monitoring mechanism for identifying distributional patterns when traditional household surveys are unavailable or unreliable.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS 1

ABSTRACT 2

TABLES 5

ILLUSTRATIONS 6

ABBREVIATIONS 7

INTRODUCTION 8

LITERATURE REVIEW 11

2.1 The Philosophy and Normative Foundations of Inequality Measurement 11

2.2 The Axiomatic Approach to Inequality Measurement 12

2.3 Definition and Measurement of Poverty: A Process. 14

2.3.1 The Foster-Greer-Thorbecke (FGT) Class 15

2.3.2 The Sen Poverty Index and Its Extensions 17

2.4 Empirical Challenges: Why the Existing Data Fail 19

2.4.1 Distributional Properties and Heavy Tails 19

2.4.2 Top-Coding and the Truncation of the Upper Tail 20

2.4.3 Underrepresentation of the Poor and Informal Economies 20

2.5 Survey Data Limitations and the Search for Alternatives 22

2.6. Administrative Data in Inequality Research 25

2.7 Housing and Vehicles as Wealth Proxies: Theoretical Rationale 27

2.8 This Paper's Position in the Literature 28

DATA AND METHODOLOGY 30

3.1. Data and Variable Construction 30

3.2. Methodology 31

3.2.1. Inequality Measures 32

3.2.2. Foster-Greer-Thorbecke (FGT) Class 35

3.2.3. Sen Index 36

3.3 Assets vs Income Inequality regressions 37

3.4 GDP growth and poverty 38

RESULTS 40

4.1. Results from Indonesia 40

4.2. Results from the U.S. 45

4.2.1. Evidence from broad inequality indexes 46

4.3. robustness check 51

DISCUSION AND CONCLUSION 53

5.1. Practical Applicability and Policy Relevance 54

5.2. Limitations and Directions for Future Reasearch 55

5.3. Summary 57

REFERENCES 59

TABLES

Table

1. Average Inequality Levels in The Positive-Only Specification 47
2. Gini Inequality Regression 50
3. Predictive Growth Model for Income and Vehicles 51
4. Average Robustness Poverty Measures, 1999–2023 52

ILLUSTRATIONS

Figure

1. Gini Coefficient Comparison between Income and Asset Proxies, IFLS Waves..... 41
2. Upper-Tail Concentration- Log Scale..... 44
3. U.S. Poverty Robustness Comparison: Income, Vehicles, and Housing..... 48

ABBREVIATIONS

IFLS: Indonesian Family Life Survey

SCF: Survey of Consumer Finances

PSID: Panel Study of Income Dynamics

GE: General Entropy

FGT: Foster-Greer-Thorbecke Poverty Class

MLD: Mean Log Deviation

DHS: Demographic and Health Surveys

SDG: Sustainable Development Goals

CHAPTER 1

INTRODUCTION

Economic inequality is one of the central preoccupations of contemporary social science and development policy. Its consequences are multidimensional and well-documented: high and persistent inequality constrains long-run economic growth, distorts labor markets, worsens health outcomes, and undermines the intergenerational transmission of opportunity. These effects are not confined to extreme cases, but they operate across a wide range of institutional settings and income levels, and they interact with poverty in ways that make the two problems difficult to address independently. A society can reduce aggregate inequality while poverty remains entrenched at the bottom of the distribution or concentrate growth among the poor without narrowing the overall dispersion of welfare. Understanding both phenomena requires not only sound conceptual frameworks but reliable empirical measurement.

That measurement is harder than it appears. Inequality and poverty are not directly observable — they must be inferred from data, and the quality of available data varies enormously across countries and over time. In high-income countries with well-functioning statistical systems, household surveys, tax records, and social insurance registers together provide a reasonably comprehensive picture of the income and wealth distribution, updated frequently enough to track change. In low- and middle-income countries, the picture is far less complete. Household surveys are conducted infrequently, cover small samples, and systematically miss both ends of the distribution: the very rich are underrepresented through non-response and underreporting, while the very poor are often excluded from sampling frames altogether. The result is that standard inequality and

poverty measures (derived almost entirely from household surveys in these contexts) are known to be biased, but the direction and magnitude of the bias are difficult to quantify.

This measurement problem has attracted considerable attention in recent years, generating a substantial literature on alternative data sources and estimation methods. Researchers have proposed asset-based welfare indices constructed from durable goods ownership, nighttime light intensity derived from satellite imagery, expenditure predictions from machine learning models applied to daytime satellite images, and wealth estimates inferred from mobile phone call records. Each of these approaches has produced genuine advances in the ability to measure welfare in data-scarce environments. Each also carries significant limitations — most critically, a dependence on conventional household survey data for model calibration that constrains their independence and limits their usefulness precisely in the contexts where surveys are most deficient.

As a necessary first step, this paper validates the distributional relationship between asset and income inequality using survey-reported asset values from the PSID, SCF, and IFLS, which serve as proxies for what administrative registries would capture. These registries exist in virtually every country, are updated as transactions occur rather than on a fixed survey schedule, and record asset values directly rather than relying on self-reported household responses. If the distribution of registered housing and vehicle values is sufficiently correlated with the distribution of income or wealth — as measured by standard inequality and poverty indices — then these records constitute a viable, low-cost, and continuously available proxy for distributional outcomes.

The key empirical question is whether this correlation is strong enough, and stable enough across institutional contexts, to support that use. We assess it using three datasets:

the Panel Survey of Income Data (PSID) covering 12 different waves from 1999 till 2022, the Survey of Consumer Finances (SCF) for the United States, covering waves from 1989 to 2022, and the Indonesian Family Life Survey (IFLS), combining three survey rounds. These settings differ substantially in income level, institutional quality, and the composition of household wealth portfolios, which allows us to assess whether the proxy relationship is robust to context or sensitive to the conditions under which it is estimated. We evaluate the relationship against a comprehensive battery of inequality and poverty measures : the Gini coefficient, the Theil index and GE family, the Foster-Greer-Thorbecke class at all three aversion levels, and the Sen poverty index, so that the assessment is not dependent on the choice of any single measure.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on inequality measurement, poverty indices, the limitations of existing data sources, and the prior use of administrative records in distributional research. Section 3 describes the survey data, and the construction of the variables used in the analysis and outlines the empirical strategy. Section 4 presents the results. Section 5 discusses and concludes.

CHAPTER 2

LITERATURE REVIEW

2.1 The Philosophy and Normative Foundations of Inequality Measurement

The measurement of inequality is inherently normative. Sen's (1973) seminal "inequality of what" question established that any measure requires a preliminary judgment about which aspect of human existence the comparison will take place within income flow, consumption expenditure, stock wealth, capabilities, or opportunity. This choice has practical consequences: different metrics produce different rankings across space and time, since these dimensions of welfare are not perfectly correlated.

The two main philosophical traditions informing distributive justice are the Rawlsian concern with the welfare of the worst-off, which focuses attention on the bottom of the distribution, and utilitarianism, which is indifferent to position and recommends attention to the whole distribution. From a practical perspective, these translate into choices about whether to use pre- or post-tax income, how to treat owner-occupied housing, and how to distinguish transitory from permanent disparities.

In practice, measuring welfare through income, consumption expenditure, or wealth has very different implications. Income captures annual flows but ignores stock wealth. Consumption is preferable in developing economies where informal incomes are hard to measure. Wealth inequality (the distribution of net assets) is more concentrated than income inequality and has received growing attention following Piketty and Saez (2003) and the World Inequality Database.

For this paper, it is important to establish whether proxy variables like housing, vehicle assets, and other buildings are more closely related to income flows or

accumulated wealth. The answer is context-specific and is addressed in the empirical sections. What matters here is that using any proxy implicitly takes a normative position: proxy validity must therefore be evaluated not merely in terms of statistical correlation, but in terms of the normative match between the proxy and the target concept.

2.2 The Axiomatic Approach to Inequality Measurement

The Axiomatic Framework for Measuring Inequality as proposed by Cowell (1980,2000) Shorrocks(1980), Bourguignon(1979), provides a different approach than the earlier mentioned philosophical approach to measuring inequality. Instead of attempting to identify how to measure inequality, so what mathematical properties should be possessed by a good index. The mathematical properties or axioms that are proposed by Cowell would provide formal definitions for commonly accepted ideas concerning how inequality should behave.

The most basic axiom is the Pigou-Dalton transfer principle, which states that a transfer of income or wealth from a higher income/wealth person to a lower income/wealth person (that does not change the relative ranking of the two), should result in a decrease in the overall measurement of inequality. If an inequality index was to increase after such a progressive redistribution, it would be considered a perverse index because it would record a higher level of Inequality even after the transfer of income/wealth from one person to another. Many easily understood summary statistics that are appropriate indices for measuring Inequality violate the principle of Pigou-Dalton transfer. For example, ratios of moments in the distribution.

Second, anonymity (or symmetry): an index measuring inequality should be based solely on the distribution of welfare levels among individuals within the population, rather than the identity of individuals. Third, population replication invariance (or replication principle): if the entire population were to be duplicated through the cloning of everyone, the index would not change due to the absence of any newly created forms of inequality; therefore, the index's relationship to any group's relative position would be maintained. Fourth, scale invariance the index is unchanged when all incomes are multiplied by the same positive constant. The most important of these axioms for empirical researchers is called additive decomposability. According to Shorrocks (1980), a decomposable index is a weighted sum of the within-group and between-group components of inequality (where groups are defined by observable characteristics such as region, educational attainment and occupation). This property is extremely useful for policymakers, as it enables researchers to identify the source of aggregate levels of inequality. Entropy-based indices such as the Theil family of indices and the squared coefficient of variation do satisfy this criterion.

The Gini coefficient however, known as the most widely used measure of inequality, fails the additive decomposability criterion. Introduced first in 1912 as a measure of economic despertion to later becoming one of the most widely used indicators of income and wealth inequality, the literature has long emphasized that it should not be interpreted as a complete or neutral representation of distributional disparities. Cowell argues that inequality measurement is not only a statistical exercise, but also depends on the way researchers conceptualize inequality and the part of the distribution they want to emphasize. This is important because different indices can produce different interpretations of the same distribution. As Cowell notes, inequality comparisons are not

always straightforward, especially when changes occur differently across the distribution: "If the rich just get richer and the poor get poorer, the answer might seem easy," but most real-world comparisons are more complex. The Gini index also does come with a number of restrictions. First off, it has maximum sensitivity towards transfers made in the middle of the income distribution, as compared to minimum sensitivity towards transfers made at either extreme of the distribution. If, for example, one compares incomes from two different countries or regions and there is an overlap in the distributions for the two groups, the comparison is not straightforward. The effects of using negative income values when considering wealth rather than income are not well-understood.

Other common indices, such as quintile and decile ratios, while clear in terms of concept and description, violate the transfer principle and provide no information about within-group distributions. Theoretical limitations of commonly used indices are not just theoretical limitations, but also limit how researchers interpret their findings and policy recommendations that arise because of their research.

For our study, we employ many indices, including the Gini index, the Theil coefficient, and the poverty indices to evaluate our results to assess the robustness of our findings to the index used. This is standard practice in the literature measuring inequality and enables us to describe the proxy relationship across various segments of the distribution.

2.3 Definition and Measurement of Poverty: A Process.

Like measuring inequality, measuring poverty has different conceptual and empirical issues. Poverty affects only those below a certain standard while equality is about all income brackets; therefore, it is logical to have different ways of computing. Multiple questions come to mind concerning the definition of a threshold (regarding

income and consumption) and how we define the "poor" from other income brackets. The World Bank's widely used method of monitoring global poverty uses consumption-based poverty lines and is regularly updated to account for changes in purchasing power parity (Ferreira et al., 2016; Jolliffe and Prydz', 2016). The definitions of poverty and consumption dramatically affect everything from the measurement of poverty's size to how it is dispersed around the world.

The research has settled on two primary ways of measuring poverty: the FGT (Foster-Greer-Thorbecke) approach to measuring poverty and the Sen index (based on Amartya Sen's work). Both measurement systems have similar axiomatic properties and are manageable for empirical research and government reports about poverty.

2.3.1 The Foster-Greer-Thorbecke (FGT) Class

Prior to Foster, Greer, and Thorbecke (1984), poverty measurement relied on a patchwork of ad hoc indices whose properties were poorly understood. The headcount ratio dominated policy reporting despite being criticized by Sen (1976) for two fundamental failures: it violates monotonicity — it does not respond to income gains among the poor that fall short of the poverty line — and it violates the transfer axiom — it is unmoved by redistributions among the poor. The income gap ratio captured the depth of poverty but remained blind to how that deprivation was distributed among the poor. Kakwani (1980) and Clark, Hemming, and Ulph (1981) proposed alternatives weighted by rank or utility, but none achieved the combination of axiomatic tractability, parametric flexibility, and decomposability that applied work required.

Foster, Greer, and Thorbecke (1984) resolved these competing demands with a single parametric family in which the poverty aversion parameter alpha governs how

much weight is given to the most deprived. At alpha equal to zero, the index reduces to the headcount ratio — the fraction of the population below the poverty line. Ravallion and Chen (1997), in their influential study of poverty trends across the developing world, relied primarily on this measure because it was the one statistic that household surveys in low-income countries could reliably support. But as Ravallion (1994) documented carefully, the headcount treats a household just below the line identically to one in extreme destitution, which means that interventions targeting the near-poor — raising them just above the threshold — register as poverty reduction even if the most severely deprived are untouched.

At alpha equal to one, the index becomes the poverty gap, averaging the normalized income shortfalls across the full population. This captures both incidence and depth and has a practical interpretation: it provides a lower-bound estimate of the aggregate resource transfer needed to close the poverty gap under perfect targeting. Fiszbein and Schady (2009), in their cross-country evaluation of conditional cash transfer programmes in Latin America, used the poverty gap to assess whether programme transfers were sufficient to close beneficiaries, consumption shortfalls. They found significant residual gaps even in well-designed programmes, illustrating the measure's value in evaluating the adequacy of redistribution. The poverty gap satisfies monotonicity but not the transfer axiom: two distributions with identical headcounts and aggregate gaps but different within-poor inequality yields identical values.

At alpha equal to two, the index squares the normalised shortfalls, assigning disproportionately greater weight to those furthest below the line. This is the only member of the FGT family that satisfies both monotonicity and the transfer axiom, making it the theoretically preferred measure when policy concern is focused on the most severely

deprived. Datt and Ravallion (1992), decomposing poverty changes in India and Brazil into growth and distributional components using all three FGT measures, showed that the choice of alpha materially affects conclusions about which source of poverty reduction dominates. Policies that produce distributionally neutral growth tend to perform well on the headcount and poverty gap but may leave the squared gap largely unchanged if growth does not reach the extreme poor.

The additive decomposability of the FGT class has made it the natural instrument for identifying geographic and demographic sources of aggregate poverty. Huppi and Ravallion (1991) used this property to decompose changes in Indonesian poverty over the 1980s into within-sector and between-sector components, finding that agricultural growth accounted for the bulk of poverty reduction but that rural-urban migration contributed an additional compositional effect. The World Bank's systematic use of FGT measures across the PovcalNet database — now the Poverty and Inequality Platform — enabled the cross-country comparisons underpinning both the Millennium Development Goals and the Sustainable Development Goals (Chen and Ravallion, 2010). Researchers routinely report all three members of the family to test whether poverty rankings are robust to the choice of aversion parameter, since rankings that hold across all values of alpha imply stochastic dominance relations between distributions (Atkinson, 1987).

2.3.2 The Sen Poverty Index and Its Extensions

Sen (1976) critique of existing poverty measures preceded and motivated the FGT contribution. His core argument was that the headcount and the income gap ratio each discard morally relevant information: the headcount is indifferent to how poor the poor are, while the income gap ratio is indifferent to how poverty is distributed among them.

Through an axiomatic derivation — identifying conditions including the focus axiom, monotonicity, and a relative equity axiom governing transfers among the poor — Sen showed that these properties uniquely characterize a composite index combining the headcount ratio, the average income gap, and the Gini coefficient of the income distribution among the poor. The index therefore rises not only when more people become poor or when the average shortfall deepens, but also when inequality among the poor increases — a dimension that the FGT class at alpha equal to one ignores entirely.

In practice, the Sen index proved difficult to apply. Its dependence on the Gini coefficient of the poverty distribution introduced sensitivity to data quality in the lower tail, and its behavior under subgroup decompositions was not well-characterized. Kakwani (1980) and Thon (1979) formalized these limitations, with Thon showing that the index violates a strong version of the transfer axiom for certain distributions. Empirical applications have nonetheless been widespread: Preston (1995) applied the index to British poverty trends across the 1980s, and Osberg and Xu (2000) used Sen-type measures for cross-country comparisons in OECD economies, arguing that the index better captures the social welfare cost of poverty than the headcount alone.

Shorrocks (1995) addressed the consistency problems by proposing the Sen-Shorrocks-Thon index, which reweights poverty gaps by each poor individual rank in the overall distribution rather than in the poverty distribution alone. This modification satisfies population replication invariance and a stronger version of the transfer axiom and decomposes into three intuitively meaningful components: the headcount, the average poverty gap among the poor, and a term capturing within-poor inequality. Osberg and Xu (2000) showed that the SST index produces poverty rankings meaningfully different from those implied by the headcount alone, particularly for countries where high

incidence of moderate poverty coexists with a long lower tail of extreme deprivation — a pattern directly relevant to the use of proxy measures, which may accurately capture incidence while distorting the lower tail.

2.4 Empirical Challenges: Why the Existing Data Fail

Even within a fixed conceptual and axiomatic framework, the empirical measurement of inequality confronts severe practical difficulties. Three deserve extended discussion here: the statistical properties of welfare distributions, the truncation of the upper tail, and the underrepresentation of the very poor.

2.4.1 Distributional Properties and Heavy Tails

Income and wealth distributions exhibit heavy upper tails that follow a Pareto distribution above a certain threshold. Inequality measures are directly affected: the Gini coefficient and variance-based measures are highly sensitive to the shape of the upper tail, so minor variations in the number or income of the very rich produce large swings in aggregate indices. Wealth distributions are more extreme still; in several developed countries, the top 1 percent holds more than 50 percent of net wealth (Saez and Zucman, 2016). The presence of negative net worths creates additional complications: the Gini is not defined for distributions with negative support, and the geometric mean is undefined for non-positive values. Cowell and Flachaire (2007) address these complications and offer robustness checks based on functional trimming and alternative functional forms.

2.4.2 Top-Coding and the Truncation of the Upper Tail

Most public-use microdata replace incomes above a set threshold with that threshold intended to protect the confidentiality of high-income respondents. Top-coding compresses the upper tail and biases Gini and related measures downward. Differential nonresponse compounds the problem: wealthy households are less likely to participate in surveys and, when they do, tend to understate assets. Atkinson, Piketty, and Saez (2011) showed that top income shares derived from tax records consistently exceed those from household surveys, with the gap widening over time as incomes at the very top have grown fastest relative to the mean.

One partial remedy is linking survey microdata to administrative tax records to impute the upper tail (Alvaredo et al.). This approach is feasible only in countries with comprehensive, accessible tax administration — a condition that excludes much of the developing world. Unlike surveys, administrative asset records of the kind proposed here do not contain intentional top-coding, because property and vehicle registration systems record transaction values rather than self-reported estimates.

2.4.3 Underrepresentation of the Poor and Informal Economies

In underdeveloped countries, many economic activities take place outside of formal institutions. For example, agriculture for subsistence (food production) and small-scale business such as grocery stores, (called petty trading) provide a livelihood for many low-income households. As a result, employment in many of these activities will not be formally recorded in wage functions or in household accounts or enterprise surveys. Hence, household surveys provide an alternative means of estimating how much money households engage in by spending on these types of goods and services.

Most survey instruments have some degree of variability due to how well enumerators administer the instruments or how well they are designed. Households at the very bottom of the income distribution (the most ultra-poor, those that are homeless, those that have been displaced or, those living in institutions) are frequently omitted from survey sampling frames, which creates a left-censoring (lower) problem, as opposed to the top-coding (upper) problem discussed earlier. Since households at the extreme ends of the income distribution are systematically excluded from the survey sampling frame, the size and structure of the total distribution of income, and thus, the estimates of inequality among low-income households, will not be accurately estimated from the measured data. On the other hand, administrative records will typically contain all households in the database and thus, they may be able to verify individuals at the lower end of the income distribution as to their assets (or properties) more accurately than households or enterprise data. However, there are different reasons for missing low-income households in each of the two administrative record sets.

Another way of estimating welfare relies on geospatial and remote sensing data through the analysis of nighttime light intensity measured from satellite images (Henderson et al., 2012). This has been shown to correlate with levels of economic output, both nationally and sub-nationally, providing a reliable estimate of welfare in places where survey data are lacking or unreliable. Jean et al. (2016) expanded upon this idea by using convolutional neural networks to process daytime satellite images to predict expenditures in five Sub-Saharan African countries, with a level of accuracy comparable to that of direct surveys. Yeh et al. (2020) demonstrated that similar techniques can be used to track changes in welfare over time at a very granular level, and Chi et al. (2022) employed machine learning to build high-resolution poverty maps

using satellite-derived features for the entire low- and middle-income world. All of these efforts have pushed the boundaries of what is possible in the field of poverty mapping and have had a significant impact on the operations of many agencies who work in international development.

2.5 Survey Data Limitations and the Search for Alternatives

The restrictions mentioned above are not simply causing statistical breakdowns, but they have led to an intensively systematic effort to develop alternative estimates of welfare and inequality. There are many directions to go for alternative approaches; the following paragraphs describe several of those alternatives, along with their benefits and limits.

The first of these alternative approaches is the longest-established route; it substitutes a currency-based direct measure of income with an asset-based indirect measure of welfare. Filmer and Pritchett (2001) proposed that the Index of Household Wellbeing should be created from the results of a principal components analysis of the ownership of household assets, including durable goods, construction materials used to build their home, and access to services. The modified Index of Household Wellbeing has been used extensively in Demographic and Health Surveys (DHS) and has undergone extensive validation and refinement (Vyas, Kumaranayake, 2006; Rutstein, Johnson, 2004). The most significant finding was that assets are a more accurate measure that represents a household's accumulated wealth and absorptive capability than flows of reported income, which are subject to recall errors and seasonal fluctuations.

The utility of asset indices as indicators of wellbeing is based on the relationship between asset ownership and the underlying wellbeing dimension of interest to us; this relationship is usually fairly strong when comparing two or more individuals from the same country at the same time, but can deteriorate significantly when comparing a single individual over time (due to the fact that asset stocks adjust to changes in income flows at a much slower rate than it does in changing income) or when comparing individuals across different countries (i.e., the same asset may mean different levels of wellbeing in different countries). These limitations are relevant for understanding the use of administrative records of assets in measuring inequality; therefore we will address these issues directly in our empirical analysis.

A second tradition uses geographical and remote sensing data to estimate wellbeing on a local scale. Many studies have shown that nighttime light as measured by satellite imagery has a positive correlation with economic activity and GDP at the country level and in smaller areas such as cities and towns (Henderson et al., 2012), and that they have been useful for estimating subnational levels of wellbeing in areas where there are few or unreliable survey data. Jean et al. (2016) extended this research by using convolutional neural networks applied to images taken during daytime from satellites to estimate how much people spend on goods and services at the village level in five Sub-Saharan African countries, achieving accuracy levels comparable to direct survey-based expenditure estimates.

Further research has refined and expanded these techniques in various ways. The study conducted by Yeh et al. (2020) examined the impact that satellite imagery could have on better predicting poverty in more specific geographic areas than would typically be possible, and this also allows researchers to track how welfare changes

through time. Chi et al. (2022) used machine learning and satellite-derived features to provide high-resolution poverty maps for all low- and middle-income countries. In addition to each of these innovations providing an unprecedented advancement in the ability to measure welfare in low- and middle-income countries, they are also driving new knowledge developments in both the academic literature as well as in the practice of international development organizations.

Digital trace data refers to the records generated by individuals as they engage in various activities (e.g., mobile phone usage, financial transactions, and online behavior). Blumenstock et al. (2015) used records of individuals' mobile phone call usage in Rwanda to predict an individual's wealth with a high degree of accuracy, and then generated a national poverty map of Rwanda at a fraction of the cost of a household survey. Kreindler et al. (2021) and Fatehkia et al. (2018) built on this methodology to create similar national poverty maps using individual-level transaction records from mobile money systems and Facebook usage data. Luo et al. (2021) also contributed to the national poverty mapping effort by using social media network structure.

The methods discussed here represent a significant advance in measuring welfare in areas with limited data, but they all share the same limitation of requiring calibration to conventional survey data at some stage. To train or validate a prediction model of welfare using input data from satellite images or phone calls, there must be labelled instances (i.e. instances that have both satellite/phone data and a measure of welfare) to create a prediction model for welfare. This means that the methods reduce, but don't eliminate, reliance on household survey data. As household survey data becomes more out-of-date and/or less frequently available, the accuracy of the predictions created by

the model decreases. Additionally, when there are no recent household surveys, the method and its associated benefits are unavailable.

To train a predictive model that relates welfare to alternative forms of data, either satellite images or call log records must have been matched to traditional household survey data; both alternative and traditional data must be available for each item used to train the model. One way to decrease dependence on household survey data is to use extrapolation technologies to derive welfare estimates using satellite imagery or call records across time and distance from traditional sources of data to predict welfare. At the point of zero recent household survey data, the extrapolation technologies offered by satellite imagery or call records will not produce a model. Thus, alternative approaches to using satellite imagery to estimate welfare are not going to rid the use of household surveys completely; however, there is a need for an alternative approach that does not depend on household survey data to calibrate a model, namely, using administrative records from government registration agencies. The only caveat to this approach is that a relationship must exist between the distribution of registered assets and welfare as defined by household surveys.

2.6. Administrative Data in Inequality Research

In the past few decades, researchers in high-income countries have studied income inequality using administrative data (records created by tax authorities, social insurance systems, land registries, and vehicle licensing agencies). Piketty (2001) used French income tax record data and Piketty and Saez (2003) utilized American income tax record data and showed that administrative data could provide the most comprehensive and accurate long-term income distribution estimates relative to household survey data. Extending their research to over 100 countries using the World

Inequality Database allowed researchers to determine that tax records met all criteria for being good measures of top income shares: covering all taxpayers, being continuously updated, and being free of recall bias and sampling errors.

Administrative data has also been used similarly to measure inequities in wages, intergenerational mobility, and effects of labour market restructuring. For example, Chetty et al. (2014) used linked U.S. earnings records to tax data from the U.S. tax system, producing a level of precision in demonstrating trends in intergenerational mobility that would not have been possible with household survey data. Card, Heining, and Kline (2013) used German social insurance records to decompose the growing wage inequality into components produced by both workers and firms. Finally, Kleven et al. (2011) used Danish administrative records to examine the relationship between tax enforcement and compliance. Across these contributions, the common theme is that administrative records offer scale, frequency, and coverage that surveys cannot match.

There hasn't been much development yet of administrative data as an alternative measure for wealth to that of income. For example, Fagereng et al. (2020) and Martínez-Toledano (2022) examined housing wealth inequality in Norway and Spain, respectively, using cadastral records and tax authority information, but both are relatively high-income countries with established administrative data systems. Vehicle registration records have only been used map of relevant areas' wealth levels indirectly, including a few studies that have validated estimates of income and wealth levels of small areas against satellite datasets, using vehicle registration records as validation instruments. To our knowledge, no previous published research has systematically examined property or vehicle registration records as stand-alone proxies of income or

wealth inequality at the country or national level and no prior research has formally tested the empirical relationship between corresponding proxies across institutional contexts. Consequently, this is the research gap addressed in the present manuscript.

2.7 Housing and Vehicles as Wealth Proxies: Theoretical Rationale

Housing typically represents the largest single asset in a household's portfolio, and vehicles constitute a substantial share of wealth for low- to moderately wealthy households. Both are durable goods whose acquisition is closely tied to savings levels and long-run earnings. The market value of a home correlates consistently with household income across a wide range of institutional settings, and vehicle and housing quality indicators are already used near-universally as welfare proxies in survey-based poverty assessments.

Administrative records of housing and vehicle transactions capture asset ownership and value — including housing size, type, condition, and location, and vehicle make, model, and age — and their aggregation across a geographic area yields a distribution of asset values whose variance and concentration reflect the underlying distribution of household wealth.

Several limitations apply. Housing and vehicles are stock variables representing accumulated wealth rather than current income and will respond slowly to economic shocks. Local asset prices reflect amenity quality as well as welfare levels. Renters, those in informal settlements, and households whose wealth is primarily financial or in livestock will not appear in these records. The proxy therefore cannot serve as a generalized welfare measure; it functions as a distributional indicator of household wealth represented by registered assets. The empirical sections examine the strength and stability

of the proxy relationship across different institutional contexts and address these limitations directly.

2.8 This Paper's Position in the Literature

The literature reviewed above defines what a valid administrative proxy must accomplish. The normative and axiomatic foundations require that it track welfare in a normatively appropriate dimension and satisfy the basic properties such as sensitivity to transfers, decomposability, that make an inequality measure useful for policy. The poverty measurement literature adds a further standard: the proxy must reproduce not only the headcount of the poor but also the depth and severity of poverty as captured by the FGT class, and the distributional structure captured by the Sen index.

Existing approaches fall short in identifiable ways. Survey data are top-coded, under-sampled at both extremes, and produced infrequently. Alternative approaches such as asset indices, satellite imagery, mobile phone data, have extended the measurement frontier but all require survey calibration, limiting their independence. Administrative tax and social insurance records have demonstrated superiority over surveys in high-income settings, but this tradition has not been extended to property and vehicle registration records as stand-alone proxies, nor applied systematically in low- and middle-income contexts.

This paper evaluates property and vehicle registration records as direct distributional proxies, using data from the United States which is a a high-income country with high-quality household surveys, and Indonesia which a lower-income country with different survey quality and wealth portfolio composition. If the proxy relationship holds

across both settings and across the full battery of inequality and poverty measures, it constitutes a viable, continuously available, low-cost alternative to infrequent household surveys for monitoring distributional outcomes.

CHAPTER 3

DATA AND METHODOLOGY

3.1. Data and Variable Construction

This study draws on data from the Panel Study of Income Dynamics (PSID), a longitudinal household survey administered by the University of Michigan. As the longest-running panel survey of its kind in the United States, the PSID tracks families and their descendants over time, collecting detailed information on income, wealth, employment, and demographic characteristics. Using this dataset, we compute measures of income and asset inequality and OLS regressions at the state level, leveraging respondents' state of residence to construct state-specific distributional estimates
The SCF is a survey that is repeated every three years by the Federal Reserve and the Treasury. Its purpose is to create a detailed database of family assets, pensions, income, and additional financial attributes for US families. High-asset households are over-sampled through a list sample using tax record data and through standard area sampling methods as well. The analysis in this research will use SCF waves from 1989 through 2022 to understand long-term structural trends in the proportional distribution of household wealth. Specifically, the wealth factors analyzed will be housing equity, auto value, non-residential structure value, total net worth, and income, both analyzed as independent entities and in various combinations.

The Indonesian Family Life Survey (IFLS) investigates households over many years, measuring many different aspects: their economy, as well as health and demographics. This research uses many different years of IFLS data (combining IFLS3, IFLS4 and IFLS5) to produce an analysis of asset holdings in total, as well as the

imputed value of rental income from these assets. The major grouping factors are urban-based vs rural-based for the Gini decomposition; and gender of the household head for the FGT decomposition. The asset groups are the same as those created for the SCF analysis (i.e. 4 groupings: housing, vehicles, non-residential buildings, and combination of these groups, total asset to the total household wealth and total household income).

3.2. Methodology

This study evaluates whether asset-based measures of inequality and poverty constitute valid proxies for their income- and wealth-based counterparts. The use of housing and vehicles as proxies for household welfare is grounded in the Permanent Income Hypothesis. Unlike annual income flows, which are subject to seasonal volatility and transitory shocks, durable assets represent "smoothed" consumption and accumulated savings. Because these assets are "lumpy" investments requiring significant capital or credit-worthiness, their distribution reflects a household's long-term economic standing and earning potential. Consequently, the distribution of registered assets is expected to carry a genuine distributional signal that mirrors the structure of broader wealth, providing a robust monitoring mechanism even when income data is unavailable.

Let I_x and P_x denote inequality and poverty measured over assets, and I_y and P_y denote the corresponding measures over income or wealth. The central question is whether $I_x \approx I_y$ and $P_x \approx P_y$ — whish is whether the distributional signals carried by asset-based measures are sufficiently close to those of income- or wealth-based measures to justify their use in contexts where the latter are unavailable or unreliable. To assess this, we employ a suite of complementary inequality measures: the Gini coefficient, the Generalized Entropy class, the Atkinson index, and percentile ratios,

alongside a set of poverty measures encompassing the Foster-Greer-Thorbecke family and the Sen index. These measures are selected not only for their broad coverage of the distributional profile, but for their distinct axiomatic properties, which together allow a multi-dimensional evaluation of proxy validity across the full distribution, within and between population subgroups, and in a regression context.

3.2.1. Inequality Measures

Let y be a random variable (housing equity, vehicle value, building value, total wealth, or income), with y_i being the value for observation i in a sample of n observations. Let $\mu = (1/n) \sum_i y_i$ denote the sample mean. All measures below are computed separately for each variable, asset combination, wave, dataset, and sample convention.

The Gini coefficient G is the most widely used scalar measure of distributional inequality. Geometrically, it is twice the area between the Lorenz curve (which plots the cumulative share of y against the cumulative share of the population ranked by y) and the 45-degree line of perfect equality:

$$G = (1 / n^2 \mu) \sum_i \sum_j | y_i - y_j |$$

$G \in [0, 1]$. A value of 0 means all households hold equal housing equity (or vehicle value, wealth, etc.); a value approaching 1 means one household holds nearly all of it. The Gini is most sensitive to transfers occurring in the middle of the distribution.

The Gini as a measure follows the properties of anonymity, which means that the identity of households does not matter, only the values y_i; scale invariance, where multiplying

every y_i by the same constant λ leaves G unchanged, so switching from USD to thousands of USD does not alter the result. Also, replication invariance; and the Pigou-Dalton transfer principle based on the idea that a transfer of one dollar from a household with more housing equity to one with less, without reversing their rank, strictly reduces G.

The GE family is parameterized by α, which controls where in the distribution the index is most sensitive:

$$GE(\alpha) = [\alpha(\alpha-1)]^{-1} \cdot (1/n) \sum_i [\, (y_i/\mu)^{\alpha} - 1 \,], \quad \alpha \neq 0, 1$$

The two limiting cases used in this study are:

$$GE(0) = (1/n) \sum_i \ln(\mu / y_i) \quad \text{[Mean Log Deviation, MLD]}$$

$$GE(1) = (1/n) \sum_i (y_i/\mu) \ln(y_i/\mu) \quad \text{[Theil T Index]}$$

The Generalized Entropy (GE) family of indices is used to measure inequality in housing, vehicle, building, and wealth and income distributions, with each member of the family differing in its sensitivity parameter α and the part of the distribution it emphasizes. GE(0), where α = 0, is most sensitive to the bottom of the distribution and measures the log-gap between the mean asset value and each household's value, meaning households with near-zero asset holdings contribute heavily to the index, making it ideal for detecting deprivation. GE(1), where α = 1, is the Theil index and offers balanced sensitivity across the full distribution by weighting each household's log-deviation by its own share of total

assets or income, making it a standard choice for cross-group or cross-state comparisons. GE(2), where $\alpha = 2$, is most sensitive to the top of the distribution and is equal to half the squared coefficient of variation; because a household holding ten times the median asset value contributes one hundred times more than one holding the median, it is particularly well-suited to capturing concentration at the very top of the wealth or income distribution.

All three GE indices satisfy anonymity, scale invariance, replication invariance, and the Pigou-Dalton transfer principle. Their defining property in this study is exact additive subgroup decomposability: for any partition of households into mutually exclusive groups (e.g., by sex or urban/rural), GE(α) decomposes exactly into a within-group component plus a between-group component, with no residual (Section 4).

The Atkinson index embeds a normative inequality-aversion parameter $\varepsilon > 0$:

$$A(\varepsilon) = 1 - (1/\mu)\,[\,(1/n)\,\textstyle\sum_i y_i^{(1-\varepsilon)}\,]^{(1/(1-\varepsilon))}, \quad \varepsilon \neq 1$$

$$A(1) = 1 - (\textstyle\prod_i y_i)^{(1/n)} / \mu \quad [\varepsilon = 1 \text{ limiting case}]$$

Three parameterizations are computed: A(0.5), A(1), and A(1.5). Higher ε assigns greater weight to inequality at the bottom. $A(\varepsilon) \in [0, 1)$ and has a direct welfare interpretation: it is the proportion of total asset value or income that could be sacrificed — if the remainder were equally distributed — while leaving social welfare unchanged at the given aversion level ε. For example, A (1) = 0.70 for housing equity means that a 70% reduction in total

housing equity, if the remainder were equally distributed, would leave welfare unchanged.

3.2.2. Foster-Greer-Thorbecke (FGT) Class

When examining Welfare Based on Assets, the identification of the relative poverty threshold (z), which helps delineate poor households, is a necessary methodological consideration. Traditional poverty lines (eg, World Bank's $2.15/day) describe income or consumption flows and do not provide a reference in terms of physical property (durable goods) to determine asset-based poverty line satisfaction. The asset poverty line for each category is set at 50 percent of the median asset value per survey wave, following the relative poverty line convention established in comparative welfare research (Atkinson, 1998; OECD, 2008; Foster, 1998). This approach is appropriate here for two reasons. First, no absolute benchmark exists for asset-based poverty — there is no equivalent of the World Bank's $2.15/day for housing equity or vehicle values. Second, a relative threshold automatically adjusts across waves and countries, making distributional comparisons consistent even as asset price levels change over time. The 50 percent threshold specifically follows the convention used in US poverty research, as distinct from the 60 percent threshold more common in European comparisons (OECD, 2008). Sensitivity to this choice is a legitimate concern; results using a 40 percent threshold are reported in the appendix as a robustness check. raBy anchoring the thresholds at median amounts, we will allow for the FGT (Foster-Greer-Thorbecke) and Sen Indexes to be responsive to changes in the overall distribution of wealth in both the United States and Indonesia and thereby provide a

consistent approach by which to identify "asset-poor" households, who have no durable holdings, compared to similarly situated households.

let y_i be household i's value of that variable. The FGT class (Foster, Greer, and Thorbecke, 1984) is:

$$P\alpha = (1/n) \sum_\{i : y_i < z \} [(z - y_i) / z]^\alpha$$

Only households below the poverty line contribute to the sum. Three values of α are computed: FGT(0) represrnting the headcount ratio, which the fraction of households with housing equity, vehicle value, building value, wealth, or income below the threshold z. It captures incidence but is insensitive to the depth of poverty FGT(1) representing the poverty gap and emphasizing the average proportional shortfall from z across all households (non-poor contribute zero). Lastly, FGT(2) or the poverty severity measure, highlighting the average squared proportional shortfall. It gives extra weight to the most asset-deprived households.

3.2.3. Sen Index

The Sen index combines three dimensions of poverty into a single scalar. Let H = FGT(0) be the headcount ratio, I = (z − μ_poor)/z be the income-gap ratio (the normalized mean shortfall of the poor from the poverty line), and G_p be the Gini coefficient computed among the poor only:

$$S \approx H\,[\,I + (1 - I)\,G_p\,]$$

The Sen index is more demanding than FGT(0) or FGT(1) alone because it penalizes inequality within the poor. If, among households below the housing equity threshold z, the distribution is highly unequal (high G_p), then S will be correspondingly higher. This makes S sensitive to the joint distribution of incidence, depth, and inequality of deprivation.

3.3 Assets vs Income Inequality regressions

Alongside the computation of inequality and poverty measures, a series of regressions has been conducted to validate furthermore this relation. We estimate a set of pooled panel regressions using state-year observations from PSID. The dependent variables are inequality measures constructed separately for income, housing, and vehicles. In this specification, We focus on the Gini coefficient under two treatments: the income variable using the zero-to-one transformation and the asset variables using the positive-only specification. This allows the analysis to compare whether the same macroeconomic and labor-market variables explain income inequality and asset inequality in a similar way.

The baseline regression is specified as:

$$I_{st}^{k} = \beta_1 UR_{st} + \beta_2 HOURS_{st} + \beta_3 GDP_{st} + \beta_4 UNION_{st} + \varepsilon_{st}$$

where I_{st}^{k} represents the Gini coefficient for inequality measure k in state ss and year t. The dependent variable changes across models: income inequality, housing inequality, and vehicle inequality. UR_{st} refers to the unemployment rate, HOURSst

captures average hours worked, GDPst measures state-level economic output, and UNIONst represents the unionization rate. The error term is denoted by εst.

This regression analysis is not designed to determine strict causal relationships. The intent of this analysis is to evaluate whether or not income and asset inequality will react similarly to the same economic conditions for each state where both types of inequality are measured. For example, if an increase in vehicle or housing inequality is positively associated with an increase in income inequality, then this would provide some evidence that assets measure some portion of the underlying distributional patterns associated with income. Conversely, if asset and income inequality did not demonstrate the same relationships (in terms of sign, magnitude, or predictive power), then this suggests that asset inequality measures a different dimension of economic standing for households than does income.

The 43 states evaluated in this study were examined over varying lengths of time (an unbalanced panel dataset). Since the models were created using this pooled sample of data (across all state-year combinations), we do not have control for the effects of time-invariant unobserved characteristics of the states. Therefore, the results of this analysis should be viewed as descriptive associations across all state-year cases rather than true fixed effect estimates.

3.4 GDP growth and poverty

To examine the macroeconomic relevance of asset-based poverty measures, we estimate a predictive growth model using a balanced panel of 43 states over a 12-year period. The primary objective is to test whether the "distributional signal" embedded in physical assets provides similar predictive power for future economic growth as traditional income metrics.

The model is specified as a Panel Data Pooling Model with the following structure:

$$GDP_growth_{i,t} = \alpha + \beta_1 Poverty_Metric_{i,t-1} + \gamma X_{i,t} + \epsilon_{i,t}$$

I utilize the Foster-Greer-Thorbecke (FGT) indices—specifically FGT1 (Poverty Gap) and FGT2 (Poverty Severity)—as the independent variables of interest. These are lagged by one period (t−1) to mitigate endogeneity and to assess their role as leading indicators. The control vector includes the state-level unemployment rate, average hours worked, and average age to account for labor market and demographic shift

CHAPTER 4

RESULTS

4.1. Results from Indonesia

The results show that asset-based measures do not all behave equally as proxies for income inequality. Across the three IFLS waves, the closest relationship with income inequality appears mainly through vehicles and housing-based combinations, while buildings alone are less consistent. This suggests that assets can partially reflect income inequality, but their performance depends on the asset type, the inequality index used, and the wave being analyzed.

A first important pattern appears in the Gini coefficient. In 2000, income inequality and vehicle inequality are almost identical under the positive-only specification. The Gini for income is 0.822, while the Gini for vehicles is also 0.822. The housing-vehicle combination is also very close, with a Gini of 0.816. This indicates that, in the 2000 wave, vehicle ownership and the combined housing-vehicle asset proxy closely reproduce the overall dispersion observed in income. The similarity is particularly important because the Gini is sensitive to the middle of the distribution, suggesting that vehicle and housing-vehicle measures capture broad inequality patterns rather than only extreme values.

In 2007, the pattern changes slightly. Income inequality is lower, with a Gini of 0.564, while housing and buildings produce somewhat higher but still relatively close Ginis of 0.602 and 0.595, respectively. The housing-building combination also remains close, with a Gini of 0.617. Vehicles, however, become less like income in this wave, with a Gini of 0.684. This suggests that the relationship between vehicles and income is

not stable across all years. In 2007, housing and building-related assets appear to track income inequality more closely than vehicles.

By 2014, vehicles again become the closest asset-based proxy to income in terms of the Gini. Income inequality has a Gini of 0.705, while vehicle inequality is 0.677. Wealth is also relatively close at 0.646, but housing and buildings are much lower, with Ginis around 0.534 and 0.539. This suggests that by 2014, housing values became less unequal than income, while vehicle ownership continued to reflect income differences more strongly.

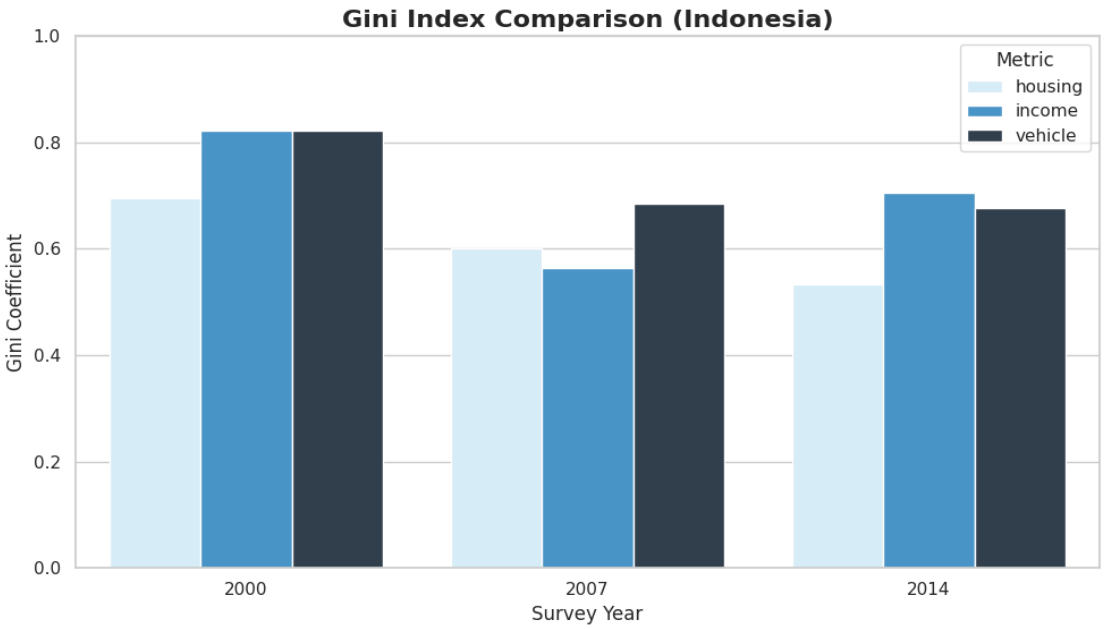

Figure 1 Gini Coefficient Comparison between Income and Asset Proxies, IFLS Waves

The Generalized Entropy indices provide a more detailed picture because they capture different parts of the distribution. GE(0), which is more sensitive to lower-tail inequality, shows that vehicles and wealth are relatively close to income in 2014. Income has a GE(0) of 1.078, while vehicles have1.012 and wealth has 1.030. This is one of the strongest results in the table because it shows that vehicles do not only match

income under the Gini, but also under a lower-tail-sensitive inequality measure. In 2007, housing and buildings are closer to income under GE(0), while in 2000 the housing-building and housing-vehicle-building combinations perform better.

However, GE (1) and especially GE(2) show that asset proxies are weaker at capturing upper-tail income inequality. This is most visible in 2014, where income GE (1) is 2.254, much higher than vehicles at 0.998, housing at 0.492, and wealth at 0.777. The difference is even larger for GE (2), while the asset measures remain far lower. This means that while vehicles and housing-related assets may reproduce the general level of inequality, they do not fully capture extreme income concentration at the top of the distribution The important conclusion drawn from this research is that there are clear signs that while wealth at the lower and middle classes appears to be represented by physical items with longevity, the absence of durable goods does not indicate that wealth is distributed evenly across the entire income distribution. Therefore, the clear shift in wealth (from physical assets that were once registered) by the ultra-wealthy (the top one percent of the income distribution) from registered physical assets to non-registered financial equities and/or to offshore accounts will remain true to this theory.

The Atkinson indices support the same interpretation. In 2000, vehicles are almost identical to income under Atkinson 0.5: income is 0.614, while vehicles are 0.615. This reinforces the Gini result and suggests that vehicles are highly consistent with income inequality when moderate inequality aversion is used. The housing-vehicle-building proxy is also close to income in 2000. In 2007, housing and buildings are again closer to income than vehicles. In 2014, vehicles remain the closest asset

proxy to income under Atkinson 0.5, although the gap becomes larger under higher inequality aversion. This shows that asset proxies become less reliable when the index places more weight on the bottom of the distribution.

The poverty-style measures, FGT (0), FGT (1), and FGT (2), also reveal an important pattern. These measures show that asset proxies do not always reproduce income poverty or deprivation in the same way they reproduce inequality. In 2000, income has an FGT (0) of 0.333, while vehicles have a higher value of 0.418 and housing has a lower value of 0.190. This suggests that housing ownership or housing values may understate deprivation relative to income, while vehicle ownership may overstate it. In 2014, however, vehicles and housing are closer to income in FGT terms. Income FGT (0) is 0.290, vehicles are 0.234, and housing is 0.261. This suggests that the ability of asset proxies to approximate income-based deprivation improves in some cases but remains uneven.

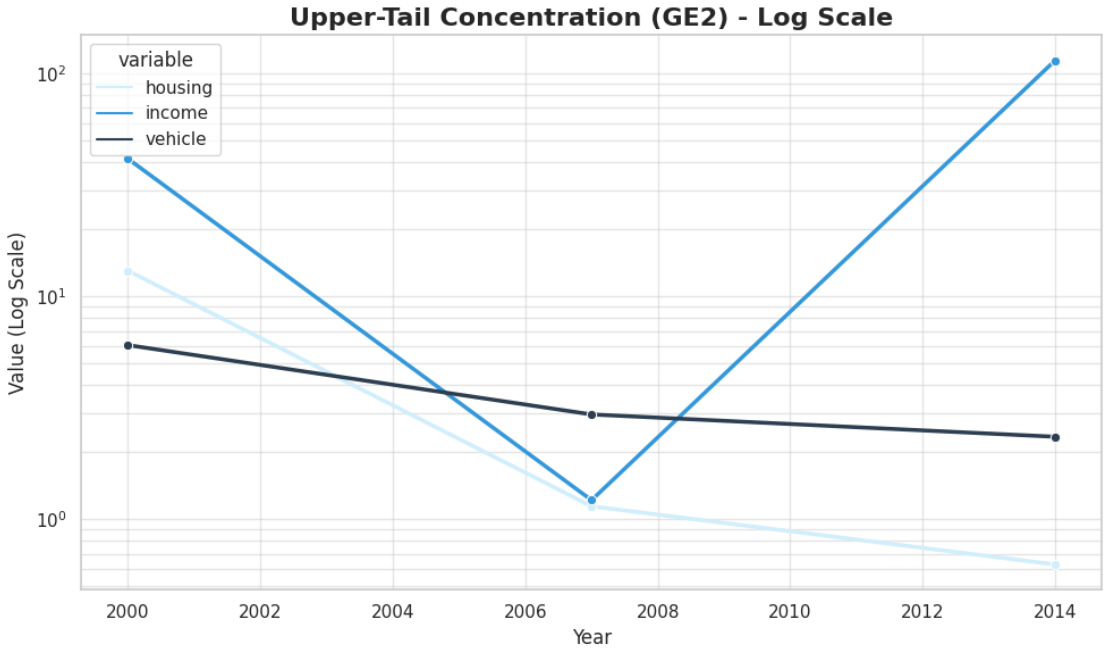

Figure 2 Upper-Tail Concentration- Log Scale

Overall, the results suggest that vehicles are the strongest single asset proxy for income inequality, especially in 2000 and 2014. Housing is also useful, particularly when combined with vehicles or buildings, but it does not consistently track income inequality across all waves. Buildings alone are the least stable proxy because their sample size is smaller, and their inequality values fluctuate more strongly across specifications.

The combined asset proxies provide additional evidence. In 2000, combinations such as housing-vehicle and housing-vehicle-building are very close to income under the Gini and Atkinson measures. This suggests that combining assets can improve proxy performance because each asset captures a different dimension of household economic position. Housing reflects long-term wealth and residence conditions, vehicles reflect liquidity and access to durable goods, and buildings capture additional real estate ownership. However, the combined proxies do not always outperform single assets. In

2014, for example, vehicles alone are closer to income than most combined measures. This means that asset aggregation improves proxy quality in some contexts, but not automatically.

The comparison between the positive-only and zeros-to-one specifications also matters. The zeros-to-one treatment produces much higher inequality for assets with many zero values, especially vehicles and buildings in 2000 and 2007. For example, building inequality under zeros-to-one reaches a Gini of about 0.964 in 2000 and 0.961 in 2007. This reflects the fact that many households do not own these assets, so replacing zeros with one creates an extremely unequal distribution. Therefore, the positive-only specification is more useful for comparing inequality among asset holders, while the zeros-to-one specification captures ownership exclusion as well as value inequality. For the main thesis story, the positive-only results are cleaner for comparing the shape of income and asset inequality, while zeros-to-one can be used as a robustness check showing how non-ownership increases measured inequality.

4.2. Results from the U.S.

The PSID results suggest that observable household assets, especially home equity and vehicles, can approximate income inequality patterns under broad distributional measures. The relationship is strongest for the Gini coefficient and Atkinson $\varepsilon = 0.5$, and it is also visible, though less cleanly, for GE(1). This provides evidence that asset-based proxies are useful when the research objective is to capture general inequality patterns rather than extreme wealth concentration.

The clearest ordering in the positive-only specification is: income inequality is lowest, home equity is slightly higher, vehicles are higher than home equity, and broader wealth measures are much more unequal. In other words, housing and vehicles

are not perfect substitutes for income, but they are much closer to income than wealth with or without equity. This supports the thesis argument that selected observable assets can act as partial proxies for income-related inequality when direct income data are unavailable, incomplete, or unreliable.

4.2.1. Evidence from broad inequality indexes

The Gini coefficient provides the strongest and most intuitive result. Across the state-year panel, the average positive-only Gini is approximately 0.431 for income, 0.477 for home equity, and 0.524 for vehicles. These values are meaningfully closer to one another than to wealth with equity and wealth without equity, which average approximately 0.710 and 0.779, respectively. This indicates that home equity and vehicles capture a moderate inequality structure, while broader wealth variables capture a much more concentrated distribution.

Atkinson $\varepsilon = 0.5$ reinforces the same conclusion. In the positive-only results, the average Atkinson 0.5 index is approximately 0.170 for income, 0.199 for home equity, and 0.235 for vehicles. Since Atkinson 0.5 is less sensitive to the bottom tail than higher Atkinson parameters, this result suggests that the similarity between income, housing, and vehicles is strongest when the index captures broad dispersion rather than extreme deprivation or concentration.

GE(1), or Theil T, also points in the same direction, especially for home equity. The average GE(1) is approximately 0.369 for income, 0.414 for home equity, and 0.506 for vehicles. This suggests that the housing proxy remains close to income under an entropy-based measure, while vehicles still move in the same direction but display somewhat higher inequality.

Table 1 Average Inequality Levels in The Positive-Only Specification

Index	Income	Home equity	Vehicles	Wealth with equity	Wealth without equity
Gini	0.431	0.477	0.524	0.710	0.779
Atkinson 0.5	0.170	0.199	0.235	0.455	0.545
GE(1)	0.369	0.414	0.506	1.113	1.425
Sen	0.137	0.164	0.173	0.295	0.273

In FGT(0), FGT(1), and FGT(2), poverty from income generally has lower levels than that from both vehicles and housing, especially during the early years. However, this difference in poverty declines in many periods of time, particularly in the ranges of 2011 - 2015 where poverty from income has risen and poverty from vehicles has decreased. This demonstrates that asset-related indicators such as vehicles do not perfectly represent poverty from income, but they can show a considerable amount of changes that occur over time in terms of household deprivation. More so, housing tends to be more stable than vehicles, with greater variations with vehicle poverty compared to housing poverty, sometimes even getting closer to income poverty. The significance of these findings for this thesis is that they illustrate how asset proxies can provide a reflection of poverty trends. However, their reliability depends upon the type of asset being measured and the type of poverty measure being used. The best reflecting measures appear to be the FGT(1) and FGT(2) measures, which are the most sensitive across all three asset types and exhibit increasing convergence for vehicles, housing, and income as the years progress, which indicates that asset proxies may provide valuable information in the assessment of the depth & severity of poverty as distinguished from indicators that measure only the number of poor people.

Figure 3 U.S. Poverty Robustness Comparison: Income, Vehicles, and Housing

Across the Atkinson specifications, the clearest pattern is the repeated similarity between income inequality and vehicle inequality. In each row, the income-based Atkinson measure and the vehicle-based Atkinson measure are estimated using either the same controls or a very similar model structure. The results suggest that vehicles are behaving in a way that is closely aligned with income inequality, especially when the dependent variables are based on the same Atkinson parameter. This is important for the thesis because it supports the idea that vehicle ownership or vehicle-based asset inequality may capture part of the same distributional structure that income inequality captures.

Another major piece of evidence supporting the above conclusion is that the pattern of similarity between the coefficients for vehicles and income in each of the

pooled models is similar, but that the broader pattern of significance of the coefficients is no less clear. For example, in the pooled model using an Atkinson of 1.5, the union co-efficient is negative (and highly statistically significant) for both the income and vehicle inequality measures (-.055 and -.054 for the union coefficient in both cases). This indicates that higher levels of unionization are related to lower levels of income and vehicle inequality. Similarly, in the fixed-effects models which were used to estimate Atkinson .5, there is also a similarly negative relationship between unemployment rates and income and vehicle inequality. Furthermore, while the significance level of the coefficient may differ marginally, the size remains close.

The first regression table on Gini coefficients shows a negative correlation between unionism and inequality across 3 measures (income, housing, vehicle). As an example, the coefficient for union is negative (-0.04897) and statistically significant ($p < 0.01$) in the income Gini model. This indicates higher unionization results in lower income inequality. The R^2 value for this model is 0.3678 based on 545 cases. Therefore, the controls account for approximately 36.8% of the variation in income inequality. Given that the dependent variable is directly income-based, this is the most powerful of the 3 models in terms of explanatory ability.

The housing and vehicle Gini regressions show a similar pattern, although with smaller R^2 values. In the housing Gini regression, the union coefficient is -0.05767* and in the vehicle Gini regression it is -0.03551*. Thus, both coefficients are negative and statistically significant, indicating there is also a negative relationship between unionism and inequality in asset-based measures (housing and vehicle). The housing Gini regression has an R^2 value of 0.0776 based on 544 cases, while the vehicle Gini regression has an R^2 value of 0.0463 based on 545 cases. Although the housing and

vehicle Gini regressions explain less variability than the income Gini regression, the sign and significance of the union coefficient remain the same across all 3 measures.

Table 2 Gini Inequality Regression

Dependent variable	Union rate coefficient	N
Income Gini	-0.04897	545
Housing Gini	-0.05767	544
Vehicle Gini	-0.03551	545

Pooled panel regressions of inequality outcomes on union rate and controls

Note: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Models are pooled panel regressions with controls for unemployment rate, average hours, and GDP.

The results of the predictive growth regressions demonstrate a high degree of consistency between income-based and asset-based poverty indicators. For both income and vehicle assets, the poverty gap exhibits a negative relationship with subsequent GDP growth. A one-unit increase in the lagged income poverty gap corresponds to a -0.107 decrease in growth, while the vehicle poverty gap shows a comparable effect of -0.116. When examining the severity of poverty, which assigns greater weight to the poorest households, the negative impact on growth intensifies. The coefficients for income and vehicle severity are -0.163 and -0.159, respectively. This suggests that extreme deprivation, regardless of the metric used to measure it, acts as a consistent drag on economic performanceNotably, the R-squared values for the income models (≈0.146) are nearly identical to those for the vehicle models (≈0.143). This indicates that the loss of information when moving from income to a physical asset proxy is minimal in a macroeconomic context.

Table 3 Predictive Growth Model for Income and Vehicles

Dependent variable	Key lagged variable	Key coefficient
GDP growth	Income FGT1, lagged	-0.10688
GDP growth	Vehicle FGT1, lagged	-0.11587
GDP growth	Income FGT2, lagged	-0.16254
GDP growth	Vehicle FGT2, lagged	-0.15897

4.3. robustness check

As a robustness check, I recompute the FGT(0), FGT(1), FGT(2), and Sen indices using a stricter poverty threshold equal to 40 percent of the median, instead of the original 50 percent median threshold. This tests whether the lower-tail patterns for income, vehicles, and housing remain similar when the poverty line is reduced.

The results remain broadly consistent with the original 50 percent threshold. Vehicles and housing continue to show higher lower-tail deprivation than income in most years, suggesting that asset-based measures capture persistent disadvantage at the bottom of the distribution. For example, in 1999, FGT(0) is 0.1506 for income, compared to 0.2412 for vehicles and 0.2129 for housing. Although the asset proxies do not perfectly replicate income poverty, the patterns remain stable under the stricter 40 percent cutoff.

This supports the robustness of the main finding that vehicles and housing can serve as useful complementary proxies for economic disadvantage when income data are limited or unreliable.

While the IFLS results test the validity of housing and vehicle proxies in a middle-income setting, the U.S. SCF results examine whether the same relationship appears in a high-income country with deeper financial markets and more complex wealth portfolios. The results show that physical asset measures, especially vehicles and combined asset proxies, continue to carry a meaningful distributional signal. Under the positive-only specification, vehicle inequality is relatively close to income inequality under broad measures such as the Gini coefficient and Atkinson 0.5. This supports the claim that observable durable assets can capture general patterns of economic stratification. However, the relationship is weaker for upper-tail-sensitive indices such as GE(2), where income and wealth concentration exceed what can be captured by housing and vehicles alone. The U.S. evidence therefore validates the proxy argument in a qualified way: asset-based measures are useful for tracking broad inequality patterns, but they should not be interpreted as full substitutes for income or wealth data, especially when the research question concerns top-end concentration

Table 4 Average Robustness Poverty Measures, 1999–2023

Variable	Average N	Average FGT(0)	Average FGT(1)	Average FGT(2)	Average Sen
Income	6,595	0.1673	0.0847	0.0574	0.1130
Vehicles	6,870	0.2218	0.1128	0.0719	0.1449
Housing	4,491	0.2228	0.1120	0.0732	0.1473

CHAPTER 5

DISCUSION AND CONCLUSION

The results reported in Chapter 4 invite a layered interpretation. Taken together, they establish that property and vehicle registration records carry a genuine distributional signal — one that is strong enough to be useful, differentiated enough to be informative, and fragile enough in specific dimensions to require careful qualification. This chapter situates those findings within the broader literature reviewed in Chapter 2, draws out their implications for the measurement of inequality and poverty in data-scarce environments, and identifies the precise boundary conditions under which the proxy relationship holds.

The main conclusion is that asset-based proxies can approximate income inequality, but only partially. The strongest evidence comes from the similarity between income and vehicles, especially in the Gini and Atkinson measures. Housing-based combinations also perform well in selected waves, especially in 2000 and 2007. However, asset proxies do not fully capture upper-tail income inequality, as shown by the large differences in GE(1) and GE(2), particularly in 2014. This suggests that assets are useful when income data are missing or unreliable, but they should be interpreted as proxies for broad economic stratification rather than exact substitutes for income.

A unified interpretation is therefore that vehicles and housing-related combinations provide the most convincing asset-based signals of income inequality, while wealth and buildings capture different dimensions of inequality. The findings support the thesis argument that asset-based measures can serve as practical proxies in data-limited settings, especially for identifying general inequality patterns. However,

they also show that no single asset fully replaces income, and that the quality of the proxy depends on whether the research question focuses on overall inequality, lower-tail deprivation, or top-end concentration.

5.1. Practical Applicability and Policy Relevance

Countries that perform household surveys only every five to ten years (the majority of low- and middle-income nations do so) are able to use administrative records (e.g., property and vehicle registration records) to provide important information about the inequality of resource distribution in-between household surveys. These Administrative records are continuously updated through the various types of transactions that occur. Thus, they provide a much more timely and cost-effective way of measuring trends in inequality than waiting for a new household survey. The evidence suggests that vehicle-based measurements of Gini and Atkinson inequality closely track income-based measurements of inequality, thereby allowing us to determine whether inequality is increasing or decreasing. When vehicle-based inequality increases, it is likely that income inequality also increases. However, vehicle registrations are less reliable indicators of the precise magnitude of change or the concentration of inequality in the upper tails of income or wealth distribution compared to income-based measures.

In addition to providing information about national trends in inequality, vehicle registration records provide a much finer level of geographic disaggregation because they are available at the individual level-of-transaction and can be aggregated based on various local government units (e.g., districts, municipalities, neighborhoods, etc.). This capability of vehicle registration records is particularly beneficial in situations where household surveys are representative only at the national level, but there is a lack of a

sample size large enough to provide reliable estimates of the incidence of poverty among subnational measures. It is theoretically possible to apply the same type of distributional analysis that the PSID used at the state level to the same type of data (i.e., administrative registration data) at much smaller spatial levels. The SDG Agenda promotes the need for an emphasis on disaggregated monitoring as well as the critical principle of leaving no one behind. In addition to providing a wealth of information about the overall level of household welfare, registration data may be especially useful for monitoring shocks to household welfare at a high frequency. While households can experience changes to their welfare very quickly as a result of crises, disasters, pandemics, or commodity price collapses, typically the timing of the surveys used to assess changes in household welfare is much later than the actual event. This is useful for vehicle-related transactions compared to other types of property-related transactions. Because vehicle transactions occur more frequently than other types of property transactions, any change in the total number (value) of transactions and/or any increase in the share of new registrations where the value is low may be a good early indicator of increased economic stress among middle-income households. While these data may require more careful interpretation than traditional measures of inequality, they provide a more timely, ecologically accurate, and policy-relevant alternative to traditional household surveys in assessing the economic disparities in a population.

5.2. Limitations and Directions for Future Research

Several limitations of the present analysis should inform the agenda for future research.

The datasets utilized in this study share a high level of longitudinal continuity as PSID/SCF and IFLS are both longitudinal datasets that were created for use to support

requisite accuracy when measuring income and wealth. Each dataset represents a "best case scenario" for validating sufficient proxy: that is: the income/wealth measures that are used to validate the asset proxies have themselves been reliably estimated and measured. The principal issue for applied use is whether the proxy relationship established here will continue to occur in locations where there are no survey data to refer to at all — the truth is, this paper does not directly test that. However, it provides evidence from a comparative analysis on the extent to which the relationship observed here persists across two very different contexts in terms of income levels and their respective institutions, which provides suggestive yet inconclusive evidence for the potential generalizability of the suggested proxy relationship.

The use of self-reported asset values obtained via household surveys as a substitute for registration record values introduces additional sources of measurement error; transaction prices are not typically recorded in the administrative registration records, but instead are recorded using self-reported values; and therefore present alternative methods for obtaining asset values entirely. There is likely to be systematic differences in the manner in which these two values operate and may exist between instances where property values are recorded for the purpose of paying taxes and where property values are recorded for property registration purposes. Future research should consider whether it can establish a direct link between household survey microdata and administrative registration records in at least one country as a means for evaluating whether the two forms of measurement exhibit congruity.

The final section will be to prove those relationships with actual administrative records from within countries where survey based, and administrative registration-based data exists. Another major area would be using countries within Latin America and/or

Southeast Asia where there are digitized property and vehicle registration systems available that could verify this study's proxy validation exercise. The implementation in these countries would take what this paper has validated by proxy and instead will directly demonstrate the proposed method in its intended use environment.

5.3. Summary

The discussion can be summarized in three propositions that capture the paper's contribution to the literature and its implications for practice:

- Asset-based proxies carry an actual distribution signal. The distributions of vehicles and housing are accurately represented by the Gini coefficients and Atkinson's measures derived from their respective distributions relative to other income-generating sources; thus creating very credible proxies for overall disparity/inequality on a broad scale within both developed (the U.S.) and developing (Indonesia) nations. This takes the Filmer/Pritchett Asset Index from a tool for ranking to providing the measure of full distributional coverage within each country's population.

- There is a theoretically established maximum limitation to the asset proxy. The asset-based proxy captures the overall median of the country-wide distribution, but systematically will underestimate upper tail concentration due to an inherent structural flaw in this design: The composition of wealth distribution changes at the upper end from registered assets to financial instruments and equities. The user of the asset proxy should see the Asset Proxy as appropriate for measuring the degree of general stratification, but not for measuring elite income concentration.

- The asset proxy can be produced independently of household surveys. While other methods (e.g., satellite imagery, mobile phone, and digital trace) require the use of a household survey to create training and/or validation estimates of prediction models, the registries of property and vehicles in the U.S. and Indonesia allow for the direct calculation of distribution statistics without the use of surveys. This lack of survey dependency is a distinct competitive advantage of the asset proxy method because many countries do not produce timely and frequent household surveys.

REFERENCES

Alvaredo, F., Atkinson, A. B., Piketty, T., and Saez, E. (2013). The top 1 percent in international and historical perspective. Journal of Economic Perspectives, 27(3), 3–20.

Atkinson, A. B. (1987). On the measurement of poverty. Econometrica, 55(4), 749–764.

Atkinson, A. B., Piketty, T., and Saez, E. (2011). Top incomes in the long run of history. Journal of Economic Literature, 49(1), 3–71.

Blanchet, T., Flores, I., and Morgan, M. (2022). The weight of the rich: Improving surveys using tax data. Journal of Economic Inequality, 20(1), 119–150.

Blumenstock, J., Cadamuro, G., and On, R. (2015). Predicting poverty and wealth from mobile phone metadata. Science, 350(6264), 1073–1076.

Bourguignon, F. (1979). Decomposable income inequality measures. Econometrica, 47(4), 901–920.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of West German wage inequality. Quarterly Journal of Economics, 128(3), 967–1015.

Chen, S., and Ravallion, M. (2010). The developing world is poorer than we thought, but no less successful in the fight against poverty. *Quarterly Journal of Economics*, 125(4), 1577–1625.

Chetty, R., Hendren, N., Kline, P., Saez, E., and Turner, N. (2014). Is the United States still a land of opportunity? Recent trends in intergenerational mobility. American Economic Review Papers and Proceedings, 104(5), 141–147.

Chi, G., Fang, H., Chatterjee, S., and Blumenstock, J. E. (2022). Microestimates of wealth for all low- and middle-income countries. Proceedings of the National Academy of Sciences, 119(3), e2113658119.

Clark, S., Hemming, R., and Ulph, D. (1981). On indices for the measurement of poverty. Economic Journal, 91(362), 515–526.

Cowell, F. A. (1980). On the structure of additive inequality measures. Review of Economic Studies, 47(3), 521–531.

Cowell, F. A. (2000). Measurement of inequality. In A. B. Atkinson and F. Bourguignon (Eds.), Handbook of Income Distribution, Vol. 1 (pp. 87–166). North-Holland: Elsevier.

Cowell, F. A., and Flachaire, E. (2007). Income distribution and inequality measurement: The problem of extreme values. Journal of Econometrics, 141(2), 1044–1072.

Datt, G., and Ravallion, M. (1992). Growth and redistribution components of changes in poverty measures: A decomposition with applications to Brazil and India in the 1980s. Journal of Development Economics, 38(2), 275–295.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2020). Heterogeneity and persistence in returns to wealth. Econometrica, 88(1), 115–170.

Fatehkia, M., Kashyap, R., and Weber, I. (2018). Using Facebook ad data to track the global digital gender gap. World Development, 107, 189–209.

Filmer, D., and 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.

Fiszbein, A., and Schady, N. (2009). Conditional Cash Transfers: Reducing Present and Future Poverty. Washington, DC: World Bank.

Henderson, J. V., Storeygard, A., and Weil, D. N. (2012). Measuring economic growth from outer space. American Economic Review, 102(2), 994–1028.

Huppi, M., and Ravallion, M. (1991). The sectoral structure of poverty during an adjustment period: Evidence for Indonesia in the mid-1980s. World Development, 19(12), 1653–1678.

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., and Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. Science, 353(6301), 790–794.

Kleven, H. J., Kreiner, C. T., and Saez, E. (2011). Why can modern governments tax so much? An agency model of firms as fiscal intermediaries. Economica, 83(330), 219–246.

Kreindler, G. E., Kreindler, A., and Buera, F. J. (2021). Measuring poverty and inequality with mobile financial data. Working Paper.

Luo, W., and colleagues. (2021). Inferring personal economic status from social network location. Nature Communications, 8, 15227.

Martínez-Toledano, C. (2022). House price cycles, wealth inequality and portfolio reshuffling. World Inequality Database Working Paper 2017/21.

Osberg, L., and Xu, K. (2000). International comparisons of poverty intensity: Index decomposition and bootstrap inference. Journal of Human Resources, 35(1), 51–81.

Piketty, T. (2001). Les hauts revenus en France au XXe siècle. Paris: Grasset.

Piketty, T., and Saez, E. (2003). Income inequality in the United States, 1913–1998. Quarterly Journal of Economics, 118(1), 1–41.

Preston, I. (1995). Sampling distributions of relative poverty statistics. Applied Statistics, 44(1), 91–99.

Pokhriyal, N., and Jacques, D. C. (2017). Combining disparate data sources for improved poverty prediction and mapping. Proceedings of the National Academy of Sciences, 114(46), E9783–E9792.

Ravallion, M., and Chen, S. (1997). What can new survey data tell us about recent changes in distribution and poverty? *World Bank Economic Review*, 11(2), 357–382.

Ravallion, M. (1994). Poverty Comparisons. Chur, Switzerland: Harwood Academic Publishers.

Rutstein, S. O., and Johnson, K. (2004). The DHS Wealth Index. DHS Comparative Reports No. 6. Calverton, MD: ORC Macro.

Saez, E., and Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. Quarterly Journal of Economics, 131(2), 519–578.

Sen, A. K. (1973). On Economic Inequality. Oxford: Clarendon Press.

Shorrocks, A. F. (1980). The class of additively decomposable inequality measures. Econometrica, 48(3), 613–625.

Thon, D. (1979). On measuring poverty. Review of Income and Wealth, 25(4), 429–439.

Vyas, S., and Kumaranayake, L. (2006). Constructing socio-economic status indices: How to use principal components analysis. Health Policy and Planning, 21(6), 459–468.

Yeh, C., Perez, A., Driscoll, A., Azzari, G., Tang, Z., Lobell, D., Ermon, S., and Burke, M. (2020). Using publicly available satellite imagery and deep learning to understand economic well-being in Africa. Nature Communications, 11, 2583.