ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

WILLINGNESS-TO-PAY: A WELFARIST REASSESSMENT

Oren Bar-Gill

Discussion Paper No. 1040

09/2020

Harvard Law School
Cambridge, MA 02138

Willingness-to-Pay: A Welfarist Reassessment

Oren Bar-Gill*

Abstract

From a welfarist perspective, Willingness to Pay (WTP) is relevant only as a proxy for individual preferences or utilities. Much of the criticism levied against the WTP criterion can be understood as saying that WTP is a bad proxy for utility – that WTP contains limited information about preferences. Specifically, the claim is that wealth effects prevent WTP from serving as a good proxy for utility. I formalize this critique and extend it. I develop a methodology for quantifying the informational content of WTP.

The informational content of WTP depends on how WTP is measured and applied. First, I distinguish between two types of policies: (i) policies that are *not* paid for by the individuals who are affected by the policy; and (ii) policies that are paid for by the individuals who are affected by the policy. Second, I distinguish between two types of WTP measures: (i) individualized WTP; and (ii) uniform, average WTP (like the VSL). When the cost of the policy is *not* borne by the affected individuals, individualized WTP has low informational content and increases wealth disparity. Uniform, average WTP has higher informational content and reduces wealth disparity, at least in the case of universal benefits. Therefore, when possible, a uniform, average WTP should be preferred in this scenario. When the cost of the policy is borne by the affected individuals, individualized WTP has high informational content but increases wealth disparity. Uniform, average WTP has lower informational content and indeterminate distributional implications. Here, the choice between individualized WTP and uniform, average WTP is more difficult.

I briefly consider two extensions. The first involves time. I present a dynamic extension of the relationship between the informational content of WTP and the wealth distribution. The second extension emphasizes the effect of forward-looking rationality on the WTP measure. The question of rationality raises additional concerns about WTP-based policymaking.

* William J. Friedman and Alicia Townsend Friedman Professor of Law and Economics, Harvard Law School. For helpful comments and discussions, I thank Matthew Adler, Jennifer Arlen, Lucian Bebchuk, Hanoch Dagan, Avihay Dorfman, Jeffrey Gordon, Duncan Kennedy, Lewis Kornhauser, Roy Kreitner, Ariel Porat, seminar participants at Tel-Aviv University, and participants at the "New Challenges for Law and Economics" conference hosted by the Center for Contract and Economic Organization at Columbia Law School and the Edmond J. Safra Center for Ethics at Tel Aviv University. I am especially grateful to Cass Sunstein for numerous, insightful exchanges on the use of willingness-to-pay in policymaking and for extensive, constructive comments on earlier drafts.

1. Introduction

Despite enduring criticism,[1] the Willingness to Pay (WTP) criterion continues to exert substantial influence on policymaking, especially through the vehicle of cost-benefit analysis.[2] Based on general Executive Orders[3] and implementing guidelines from the Office of Management and Budget (OMB),[4] rulemaking by government agencies – from the Environmental Protection Agency[5] to the Department of Transportation[6] (DOT) to the Department of Health and Human Services[7] (HHS) to the Nuclear Regulatory Commission (NRC)[8] – relies on WTP-based cost-benefit analysis.[9]

In this Article, I use a welfarist framework to evaluate WTP-based policymaking and the critiques of WTP-based policymaking. I begin with the fundamental question: Why use WTP at all? A

---

[1] See, e.g., C. Edwin Baker, *The Ideology of the Economic Analysis of Law*, 5 PHIL. & PUB. AFF. 3, 16–19 (1975); Lucian A. Bebchuk, *The Pursuit of a Bigger Pie: Can Everyone Expect a Bigger Slice?*, 8 HOFSTRA L. REV. 671 (1980); Hanoch Dagan, *Political Money*, 8 ELECTION L.J. 349, 356 (2009); Ronald M. Dworkin, *Is Wealth a Value?*, 9 J. LEGAL STUD. 191 (1980); Duncan Kennedy, *Cost-Benefit Analysis of Entitlement Problems: A Critique*, 33 STAN. L. REV. 387 (1981); Duncan Kennedy, *Law-and-Economics from the Perspective of Critical Legal Studies*, *in* 2 NEW PALGRAVE DICTIONARY OF ECONOMICS AND THE LAW 465, 471–72 (Peter Newman ed., 2002); Anthony T. Kronman, *Wealth Maximization as a Normative Principle*, 9 J. LEGAL STUD. 227, 240 (1980); Zachary Liscow, *Is Efficiency Biased?*, 85 U. CHI. L. REV. 1649–1718 (2018).

[2] For cost-benefit analysis in U.S. policymaking, see, e.g., Exec. Order No. 12,866, 3 C.F.R. 638 (1994) (requiring cost-benefit analysis in federal agencies); see also Michigan v. EPA, 135 S. Ct. 2699, 2706–07 (2015) (holding that the EPA was required to consider cost against benefits before promulgating the regulatory scheme at issue). For the influence on cost-benefit analysis on policy-oriented research, see, e.g., COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC AND PHILOSOPHICAL PERSPECTIVES (Matthew Adler & Eric A. Posner, eds., 2000); RICHARD L. REVESZ & MICHAEL A. LIVERMORE, RETAKING RATIONALITY: HOW COST-BENEFIT ANALYSIS CAN BETTER PROTECT THE ENVIRONMENT AND OUR HEALTH (2008); Cass Sunstein, *The Real World of Cost-Benefit Analysis: Third-Six Questions (And Almost As Many Answers)*, 114 COLUM. L. REV. 167 (2014).

[3] See, e.g., Exec. Order No. 12,866, 3 C.F.R. 638 (1994).

[4] OFFICE OF MGMT. & BUDGET, EXEC. OFFICE OF THE PRESIDENT, CIRCULAR A-4, 18–20 (2003) ("'Opportunity cost' is the appropriate concept for valuing both benefits and costs. The principle of 'willingness-to-pay' (WTP) captures the notion of opportunity cost by measuring what individuals are willing to forgo to enjoy a particular benefit.").

[5] U.S. ENVTL. PROT. AGENCY, GUIDELINES FOR PREPARING ECONOMIC ANALYSES: MORTALITY RISK VALUATION ESTIMATES, at B-4 (2010).

[6] U.S. DEP'T TRANSP., REVISED DEPARTMENTAL GUIDANCE 2016: TREATMENT OF THE VALUE OF PREVENTING FATALITIES AND INJURIES IN PREPARING ECONOMIC ANALYSES (2016), at 1; U.S. DEP'T TRANSP. THE VALUE OF TRAVEL TIME SAVINGS: DEPARTMENTAL GUIDANCE FOR CONDUCTING ECONOMIC EVALUATIONS REVISION 2 (2016 UPDATE), at 2.

[7] U.S. DEP'T HEALTH AND HUMAN SERVICES, VALUING TIME IN U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES REGULATORY IMPACT ANALYSES: CONCEPTUAL FRAMEWORK AND BEST PRACTICES, at 7–11 (2017) (stating that time used—a cost in HHS's cost-benefit analysis for a given set of regulation—is a function of wages, which are, in turn, a function of a worker's WTA).

[8] U.S. NUCLEAR REGULATORY COMM'N, COST-BENEFIT GUIDANCE UPDATE, at 16 (2017) ("NRC utilizes the willingness to pay (WTP) method for calculating VSL [value of a statistical life], consistent with other Federal agencies.")

[9] Many of the above examples are taken from Liscow, *supra* note 1. The WTP criterion is also invoked in tort law. See, e.g., Jennifer H. Arlen, *An Economic Analysis of Tort Damages for Wrongful Death*, 60 N.Y.U. L. Rev. 1113 (1985); Ariel Porat and Avraham Tabbach, *Willingness to Pay, Death, Wealth, and Damages*, 13 Amer. L. Econ. Rev. 45 (2011). While my focus is on regulatory decisionmaking, some of the analysis may also be relevant to tort law. In some cases, e.g., when a legal policy removes an existing entitlement, Willingness-to-Accept (WTA) may be more appropriate than WTP. The analysis in this paper would apply for policymaking based on WTA, although the WTA measure should be less sensitive to wealth.

welfarist cares about individual preferences or utilities – and how they are aggregated into a social welfare function – not about individuals' WTP. The answer is simple: We cannot directly observe preferences or utilities. And so we use WTP as a proxy for utility. The idea is that WTP contains important information about preferences and utility.[10] Much of the criticism levied against the WTP criterion can be understood as saying that WTP is a bad proxy for utility – that WTP contains limited information about preferences.[11]

1.1 The Informational Content of WTP

The main goal of this Article is to explore the conditions under which WTP can serve as a good proxy for utility. A major criticism of WTP is that wealth effects prevent WTP from serving as a good proxy for utility. I formalize this critique and extend it. In particular, I analyze the effects of the distribution of wealth in society on the informational content of WTP. The basic claim is that WTP contains more information about preferences, and thus serves as a better proxy for utility, when the distribution of wealth is more equal. Conversely, in a society with great wealth disparities, there is a greater risk that WTP will be a poor proxy for utility.[12] Whether or not this risk is realized, critically depends on how WTP is measured and applied.

I develop a methodology for quantifying the informational content of WTP. This methodology requires the specification of a functional relationship between wealth and utility, which captures the decreasing marginal utility from money. More fundamentally, this functional relationship is assumed to be common across individuals (representing common personal preferences), which supports cardinal and interpersonally comparable utilities.[13] The power of this methodology is demonstrated using a particular functional relationship that is borrowed from other applications in the economic literature and supported by data. My purpose, however, is not to defend any specific function, but rather to show how the distortion caused by WTP-based policymaking can be quantified *given* a specified functional relationship between wealth and utility.

As mentioned above, the informational content of WTP depends on how it is measured and on the policy choices that WTP is called upon to resolve. I start with a taxonomy that identifies and distinguishes the relevant scenarios.[14]

---

[10] See, e.g., ANDREU MAS-COLELL ET AL., MICROECONOMIC THEORY, Ch. 3, Sec. 3.D., p. 50 (1995) (discussing demand theory; demand curves aggregate WTP across all consumers).

[11] A note on terminology: I use the terms "preferences," "utility" and "social welfare function" as they are defined in microeconomics and welfare economics. See, e.g., MAS-COLELL ET AL., *id*, at 3-14 (on preferences and utility), 789-790 (on social welfare) (1995). A utility function represents a preference ordering, and a social welfare function is an aggregation of individuals' utility functions. Accordingly, my focus is on the preference-satisfaction version of welfarism. Although at least some of the arguments apply to other versions of welfarism. On the different versions of welfarism – see, e.g., MATTHEW D. ADLER, MEASURING SOCIAL WELFARE: AN INTRODUCTION, pp. 10-11 (2019).

[12] See, e.g., ADLER, *id*, at 35 ("CBA's valuations are skewed by the diminishing marginal well-being impact of money.").

[13] The proposed methodology is closely related to the distributional weights approach in cost-benefit analysis. For an excellent exposition to this approach – see Matthew D. Adler, Benefit-Cost Analysis and Distributional Weights: An Overview, 10 Rev. Environ. Econ. & Policy 264 (2016).

[14] Compare: Adler, supra note 13, at 275-277. Adler's distinction between different "cost incidence" is similar to my distinction between policies that are, and are not, paid-for by the individuals who are affected by the policies. Also,

1.1.1 WTP Measures and Uses: A Taxonomy

When considering the effects of WTP on policymaking, it is important to distinguish between two types of policies, based on who bears the cost of the policy:

- Policies that are *not* paid for by the individuals who are affected by the policy, but rather by general funding sources (like tax revenues). Consider policies that improve the country's schools, e.g., by hiring more teachers, improving teacher training, upgrading school buildings or the deployment of new technology. These policies are not paid for by the students or their families, but rather by general tax revenues.

- Policies that are paid for by the individuals who are affected by the policy. Consider policies that improve car safety, e.g., by mandating features like airbags, anti-lock braking systems (ABS) or rearview cameras. Regulation that mandates such features will increase the cost of manufacturing cars, and this cost will be passed on (at least in part) to car buyers. Therefore, car owners who benefit from the policy by driving safer cars also bear the cost of the policy.[15]

Another important distinction is between two types of WTP measures:

- Individualized WTP – Measures the benefit from a policy by eliciting the WTP of the individuals who are affected by the policy. Consider a policy that reduces the mortality risk of individuals in a certain geographic location (e.g., by improving air quality in the region). An individualized WTP asks the affected individuals how much they would pay for the reduction in mortality risk.

- Uniform, average WTP – Measures a universal benefit by eliciting and aggregating the WTP of all individuals. Reducing mortality risk is an example of a universal benefit. Some policies reduce mortality risk for one group of individuals, whereas other policies reduce mortality risk for a different group of individuals. Policymakers can elicit WTP for a reduction in mortality risk from all individuals, regardless of any specific policy. Policymakers can then calculate the average WTP across the population and use this average figure as a uniform WTP whenever a policy affects mortality risk. Returning to the policy that reduces mortality risk in a certain geographic location, this approach would use the uniform, average WTP, rather than elicit WTP from the individuals in that geographic location.

Individualized WTP measures, or at least WTP measures that are disaggregated by income groups, are sometimes used in practice.[16] For example, DOT uses a WTP-based measure of time, called

---

Adler's distinction between differentiated values and population-average values parallels my distinction between individualized WTP and uniform, average WTP.

[15] Policies that are and are not paid-for by the affected individuals mark two polar extremes. Between these extremes, lie many policies that are partially funded by the affected individuals. For example, students and their families may pay, at least partially, for higher-quality schools through school-attendance fees and higher property taxes. And car manufacturers may not be able to pass all of the increased cost to car buyers.

[16] *See* Liscow, *supra* note 1. *See also* Cass R. Sunstein, *Valuing Life: A Plea for Disaggregation*, 54 DUKE L.J. 385 (2004) (arguing for the use of disaggregated WTP values); ADLER, supra note 11, at 199 (noting that "textbook" cost-benefit analysis uses individualized monetary equivalents).

Value of Time Travel Savings (VTTS). The VTTS does not have a single, uniform value; rather it is higher for air and high-speed rail travel and lower for intercity travel (buses). DOT has adopted an explicitly income-based justification for the different VTTS values – users of air and high-speed rail are richer than those who ride the bus and thus would be willing to pay more time saved.[17] In most cases, however, policymakers use a uniform, average (or median) WTP, aggregated across the entire population, even when the policy affects only a subset of the population. Most prominently, the value of a statistical life (VSL) that is routinely used in the cost-benefit analysis of regulations that affect mortality risk is a uniform, population-wide figure.[18]

The preceding distinctions are summarized in the following 2x2 table. I will show that an assessment of the informational content of WTP must be sensitive to the specific cell in this table.[19]

		What type of WTP measure is used?	
		Individualized WTP	Uniform, Average WTP
Who pays for the policy?	Not those affected by the policy	Scenario I	Scenario II
	Those affected by the policy	Scenario III	Scenario IV

Table 1: Different Policy Types and Different WTP Measures

1.1.2 Policies that are *Not* Paid-for by the Affected Individuals (Scenarios I and II)

When WTP is used to evaluate policies that are not paid-for by the affected individuals, the main concern is about the informational content of individualized WTP (Scenario I). In this context, informational content can be conceptualized as follows: A WTP measure is perfectly informative, when it supports the adoption of Policy A rather than Policy B, if and only if Policy A creates

---

[17] *See* The Value of Travel Time Savings: Departmental Guidance for Conducting Economic Evaluations Revision 2 (2016 Update), U.S. DEP'T TRANSP. 7 (2016).

[18] For an excellent account of how VSL figures are derived and how they are used in policymaking – see Sunstein, *Id*, at 396-404. There are two sources for VSL figures. The first is average responses from contingent valuation studies. The second and more influential source is market evidence – from labor markets and consumers markets – on the price of safety, e.g., the wage premium for a job that entails a certain mortality risk. *Id*. The market value of mortality risk can be thought of as an average WTP figure, if market participants – in a given market or across different markets used to derive the uniform VSL – are representative of the general population. Note that market wages, for example, are affected by the risk premium demanded by the marginal employee, not by the average employee. If marginal employees and marginal consumers, across different markets, are systematically poorer, then the VSL is not an average WTP figure. When the VSL or, more generally, WTP is derived from market transactions, there might be distortions: (1) Market failures might bias attempts to derive WTP information from market prices; (2) WTP for a benefit provided through the market may be different from WTP for the same benefit provided by the government These and other distortions, in the measurement of WTP, are not addressed in this paper.

[19] The four scenarios in Table 1 are theoretical archetypes. Real-world policymaking is often a hybrid of two or more scenarios.

greater utility than Policy B. The informational content of the WTP measure goes down, when it supports Policy A, even though Policy B creates greater utility. In particular, an individualized WTP would support a Policy A that benefits a rich Individual A, even though Policy B creates greater utility for a poor Individual B. The methodology developed in this Article allows us to quantify this distortion by calculating the maximal ratio between the utility that would have been created by the rejected Policy B and the utility that is created by the adopted Policy A. Consider an illustrative example based on US data. In the example, when Individual A's wealth is in the 70th percentile and Individual B's wealth is in the 30th percentile (which means that Individual A's wealth is 14.7 times greater than Individual B's wealth), then a WTP-based assessment will support Individual A's preferred policy, Policy A, even when the benefit provided by Individual B's preferred policy, Policy B, is orders of magnitude greater, specifically, 14.29 – 200 times greater, depending on the scale of the policies considered.

In some cases, the use of uniform, average WTP figures (Scenario II), instead of individualized WTP figures (Scenario I), reduces, or even eliminates, the distortion caused by wealth disparity. Consider Policy A that benefits the rich and saves 1,000 (statistical) lives, and Policy B that benefits the poor and saves 2,000 (statistical) lives. With individualized WTP, the less effective Policy A might be preferred. But if a uniform, average WTP is used for measuring reduction in mortality risk (VSL), then the more effective Policy B will be preferred. The uniform, average WTP has greater informational content than the individualized WTP.[20] But uniform, average WTP measures do not always solve (or mitigate) the wealth-disparity problem, and they do not always increase the informational content of WTP. A uniform, average WTP is informative only when measuring a universal benefit, like reduction in mortality risk, that everyone cares about.

1.1.3 Policies that are Paid-for by the Affected Individuals (Scenarios III and IV)

When WTP is used to evaluate policies that are paid-for by the affected individuals, the main concern is about the informational content of uniform, average WTP (Scenario IV). It is useful to begin with the individualized WTP, and explain why the individualized measure has high informational content in this scenario. The poor are willing to pay less than the rich for a policy that would create the same (or greater) utility, because the poor have other, high-utility uses for the little money that they have, e.g., paying rent and buying food. The rich, on the other hand, have more money and lower-utility uses for their marginal dollars (think of a billionaire buying her 10th yacht). Individualized WTP thus balances the utility created by the policy (- the benefit side) against the utility from alternative uses (- the cost side). Since both benefits and costs are important, individualized WTP is normatively appealing. When the cost of implementing the policies is borne by the individuals who are affected by these policies, adopting a policy that affects the rich and rejecting an equal-benefit policy that affects the poor is a feature, not a bug. This outcome reflects the high informational content of individualized WTP.

When policymakers replace the individualized WTP with a uniform, average WTP, informational content might be lost. Consider a Policy B that reduces mortality risk for the poor Individual B, e.g., a regulation that forces manufacturers to add a certain safety feature to a product that is

[20] See Matthew D. Adler & Eric A. Posner, Implementing Cost-Benefit Analysis When Preferences Are Distorted, 29 J. LEGAL STUD. 1105, 1122-23 (2000) (explaining that "a constant figure for the monetized value of life" is one way to address the distortion caused by wealth disparity).

purchased mainly by the poor and that would thus increase the price of the product. While Individual B clearly benefits from the reduction in mortality risk, if this reduction is not very high Individual B might not be willing to pay the higher price for the safer product. In this case, with individualized WTP, Policy B would be rejected. But, if policymakers use the higher, average WTP for a reduction in mortality risk, then Policy B might be adopted. Or consider a Policy A that reduces mortality risk for the rich Individual A, e.g., the safety feature is now added to a product that is purchased mainly by the rich. Individual A may be happy to pay the higher price for the safer product. And thus, with individualized WTP, Policy A would be adopted. But, if policymakers use the lower, average WTP for a reduction in mortality risk, then Policy A might be rejected.

These distortions have been identified in the literature.[21] I quantify them by showing how greater wealth disparity increases the range of welfare-reducing policies that would be adopted if a uniform, average WTP is used. A larger distortion means lower informational content of the WTP measure. Extending the example described above, where Individual A's wealth is in the 70th percentile and Individual B's wealth is in the 30th percentile, I show that with a uniform, average WTP, the policymaker might adopt a welfare-reducing policy that will force Individual B to pay up to 9.38 times as much as the benefit is actually worth to him, or up to 838% more than he is willing to pay for the policy. Similarly, the policymaker might fail to adopt a policy that costs much less than what Individual A would be willing to pay for the policy.

1.1.4 Summary

When the cost of the policy is not borne by the affected individuals, uniform, average WTP has more informational content than individualized WTP. In contrast, when the cost of the policy is borne by the affected individuals, uniform, average WTP has less informational content than individualized WTP. In both scenarios, however, when WTP – either individualized WTP or average WTP – has limited informational content, this limit is increasing in the degree of wealth disparity.

These results suggest a previously underappreciated social cost of wealth disparity and present a novel challenge for WTP-based policymaking. If greater inequality reduces the informational content of the WTP measure, then a policy that exacerbates wealth disparities will make it harder to identify welfare-enhancing policies in the future. From a welfarist perspective, the main justification for using WTP is the information it carries about preferences and utility. If WTP-based policymaking exacerbates wealth disparities, then the mere use of WTP in policymaking undermines the justification of using WTP in policymaking. WTP-based policymaking might become self-defeating. A key question, therefore, is: When does WTP-based policymaking exacerbate wealth disparities?

When the cost of the policy is *not* borne by the affected individuals, greater wealth disparity reduces the informational content of individualized WTP. In this scenario, WTP distorts policymaking in a particular direction – benefiting the rich at the expense of the poor.[22] Using a

---

[21] See Sunstein, *supra* note **Error! Bookmark not defined.**.

[22] See, e.g., ADLER, *supra* note 12, at 34 (providing an example in which the wealthier group, as opposed to a poorer group, is willing to pay more for medical treatment and therefore receives that treatment from the government); Baker,

uniform, average WTP reduces, and even eliminates, the distortion in some cases – when the only relevant benefit of the considered policies is a universal benefit, like a reduction in mortality risk – but not in all cases. Moreover, in the important case of universal benefits, a uniform, average WTP can support progressive redistribution (or, at least, avoid the regressive redistribution of individualized WTP), regardless of informational content. Consider a policy that saves many (statistical) lives of poor individuals, but costs billions to implement. Using the poor individuals' WTP, the policymaker might conclude that the benefit does not justify the cost and reject the policy. Using the higher, average WTP, the same policy may be adopted. (Using the average WTP and adopting the policy is especially good for the poor, if the implementation costs are paid for by general taxes and the poor pay less taxes.) Now consider a policy that saves many (statistical) lives of rich individuals. Using the high WTP of the rich, the policy would be adopted, despite high implementation costs. The same policy may be rejected if we use the lower, average WTP.

When the cost of the policy is borne by the affected individuals, individualized WTP has high informational content. When considering policies that affect the poor, individualized WTP does not force the poor to pay more than they can for a benefit. And when considering policies that affect the rich, individualized WTP supports high-cost policies that create even higher benefits. The high informational content, however, does not prevent individualized WTP from supporting policy choices that increase wealth disparity. When the cost of the policy is borne by the affected individuals, a uniform, average WTP has lower informational content. It harms both the rich and the poor, with indeterminate distributional implications.

To summarize: When the cost of the policy is *not* borne by the affected individuals, individualized WTP has low informational content and increases wealth disparity. Uniform, average WTP has higher informational content and reduces wealth disparity, at least in the case of universal benefits. Therefore, when possible, a uniform, average WTP should be preferred in this scenario. When the cost of the policy is borne by the affected individuals, individualized WTP has high informational content but increases wealth disparity. Uniform, average WTP has lower informational content and indeterminate distributional implications. Here, the choice between individualized WTP and uniform, average WTP is more difficult. The analysis provides further justification for the common use of uniform, average WTP measures, but only when the cost of the policy is *not* borne by the affected individuals.[23]

---

supra note 1, at 9 ("As a general matter, the rich are favored . . . to the extent that the rich own a disproportionate share of the productive assets, or more strictly, to the extent that the rich are more likely to be willing and able to buy a right for productive use."); Bebchuk, *supra* note 1; Dworkin, *supra* note 1, at 199–200; Kennedy, *A Critique*, *supra* note 1; Kennedy, *Law-and-Economics*, *supra* note 1; Liscow, *supra* note 1, at 1652 ("Because the rich have greater wealth, the view goes, they will tend to have a greater willingness to pay, and therefore policymakers maximizing efficiency will choose policies that benefit the rich over the poor. . . .).

[23] While my focus is on identifying and quantifying the distortion caused by wealth disparity, the proposed analytical framework can also be used to correct for the wealth disparity – to derive a wealth-adjusted WTP that policymakers can apply. When the cost of the policy is *not* borne by the affected individuals and we have an empirically assessed individualized WTP of a poor individual or a rich individual, I show how to derive the WTP of an individual with median wealth for the same policy or benefit. The proposed wealth adjustment is closely related to the distributional weights approach. See Adler, supra note 13. When the cost of the policy is borne by the affected individuals, the problem is with the uniform, average WTP; and it can be corrected by shifting to an individualized WTP, which has higher informational content.

1.2 Time and Rationality

Two extensions are briefly considered. The first involves time. I present a dynamic extension of the relationship between the informational content of WTP and the wealth distribution. Since WTP is affected by wealth, the initial wealth distribution will affect the policies that a WTP-based analysis prescribes. But these chosen policies will then change the distribution of wealth, which will then change WTP and thus lead to further policy change. This further policy change will again affect the distribution of wealth. Etc. Through this dynamic, inequality can increase over time.[24] (Inequality can also decrease over time, if only uniform, average WTP is used.)

A standard critique of WTP-based policymaking is that the chosen policy depends on the initial distribution of wealth.[25] The dynamic extension strengthens this critique. The initial wealth distribution affects not only the current policy choice, but also many future policy choices. In addition, the dynamic extension forces us to rethink the WTP for the initial policy. Since the initial policy will affect, through the evolving wealth distribution, many future policies, the stakes are higher and thus WTP for the initial policy will be higher. Indeed, individuals would borrow against future wealth to increase WTP and secure their favored initial policy.

The second extension emphasizes the effect of forward-looking rationality on the WTP measure. Consider the standard WTP question: "how much are you willing to pay for Policy X?" For a rational individual, this question would elicit a response that is sensitive to changes in the wealth distribution brought about by the policy – in the short-term and in the long-term (incorporating the dynamic extension). A myopic individual, on the other hand, will consider only the immediate effects of the policy, ignoring its implications for the wealth distribution and for future policy debates. Therefore, the question of rationality raises additional concerns about WTP-based policymaking.[26]

***

The remainder of the Article is organized as follows. Section 2 lays out the framework of analysis. Section 3 analyzes, and quantifies, the effects of the wealth distribution on the informational content of WTP, when the cost of the policy is *not* borne by the affected individuals. Section 4 shifts the focus to scenarios, where the cost of the policy is borne by the affected individuals. The two extensions – concerning time and rationality – are briefly discussed in Section 5. Section 6 concludes.

---

[24] There are other reasons why the rich get richer. *See*, *e.g.*, DANIEL RIGNEY, THE MATTHEW EFFECT: HOW ADVANTAGE BEGETS FURTHER ADVANTAGE (2010) (suggesting that economic Matthew effects occur among other reasons, due to inheritance, compounding interest, promotion and compensation dynamics, and monopoly and oligopoly effects); Liscow, *supra* note 1.

[25] See, e.g., Bebchuk, *supra* note 1; Kennedy, STAN. L. REV., *supra* note 1.

[26] For a different critique of WTP-based policymaking that is also based on bounded rationality – see Sunstein, *supra* note **Error! Bookmark not defined.**, at 403, 411, 427-28.

2. Framework of Analysis

2.1 Setup

Consider a society with two individuals, Individual A and Individual B (or, equivalently, a society with two homogeneous groups, Group A and Group B). The parties' utilities are denoted by u_A and u_B. And their wealth is denoted by ω_A and ω_B. I assume, without loss of generality, that $\omega_A > \omega_B > 0$. Specifically, let $\omega_A = \gamma\omega_B$, with $\gamma > 1$. Let $\overline{\omega} \equiv (\omega_A, \omega_B)$ denote the vector of wealth values.[27] The social welfare function is: $W(u_A, u_B)$.

Denote the status quo policy by P^0. In the status quo, individual utilities are: u_A^0 and u_B^0, and social welfare is $W^0 = W(u_A^0, u_B^0)$. A new policy, P^1, is being considered. This new policy increases Individual A's utility by Δu_A and Individual B's utility by Δu_B. For simplicity, we assume that $\Delta u_A \geq 0$ and $\Delta u_B \geq 0$. Let $\Delta u_A = \delta \Delta u_B$, with $\delta \geq 0$.[28] The utility changes, Δu_A and Δu_B, reflect the benefits from P^1, e.g., cleaner air or better schools. If the individuals need to pay for P^1, e.g., through higher taxes or higher product prices, then these costs will also affect the individuals' utilities under P^1. We thus have $u_A^1 = u_A^0 + \Delta u_A$ and $u_B^1 = u_B^0 + \Delta u_B$. Social welfare, under P^1, will be $W^1 = W(u_A^1, u_B^1)$.

The new policy, P^1, can affect only Individual A or only Individual B. Denote by 'Policy A' a policy that affects only Individual A and denote by 'Policy B' a policy that affects only Individual B. Specifically, Policy A increases Individual A's utility by Δu_A, and Policy B increases Individual B's utility by Δu_B. The idea is to distinguish between policies that affect the rich (Individual A) and policies that affect the poor (Individual B). Of course, there are also policies that affect both the rich Individual A and the poor Individual B. But, for present purposes, it is sufficient to focus on the two targeted policies, Policy A and Policy B. The analysis can be extended to account for hybrid policies.[29]

---

[27] Utility is affected by wealth – by the individual's wealth and by the overall distribution of wealth. If I have more wealth, then (other things being equal) I can consume more and thus increase my utility. But if everyone's wealth increases, then prices may increase such that my personal increased wealth does not translate into more consumption. Also, an individual may independently care about how her own wealth compares to that of others in the population. We thus have: $u_A(\omega_A, \overline{\omega})$ and $u_B(\omega_B, \overline{\omega})$. In the examples studied in Sections 3 and Section 4 (below), we define wealth, ω_i, in relation to the average or median wealth level, rather than in absolute dollar terms, thus accounting for these relative wealth effects.

[28] The policy P^1 affects utilities through two channels: (1) The direct channel – the policy brings about a new state of the world that is better for at least some individuals, e.g., shorter wait times at airports that increase the utilities of travelers. And (2) the indirect, wealth channel – the policy changes the distribution of wealth, which then affects utilities. For example, if I can get to the airport an hour later, then I can stay at work an hour longer and earn more money, which I can then spend on consumption. (Some policies directly affect the distribution of wealth, e.g., policies that change the level of taxes or subsidies, and policies that grant monopoly power, through IPRs or otherwise.)

[29] A more general model would define a continuous range of policies, with Policy A that affects only Individual A at one end of the range and Policy B that affects only individual B at the other end of the range. Specifically, let $\phi \in [0,1]$ and define a Policy ϕ that creates utility $\Delta u_A(\phi)$ for Individual A and utility $\Delta u_B(\phi)$ for Individual B. Assume that $\Delta u_A(\phi = 0) = 0$, $\Delta u_A{}'(\phi) > 0$ and $\Delta u_A{}''(\phi) \leq 0$; and assume that $\Delta u_B(\phi = 1) = 0$, $\Delta u_B{}'(\phi) < 0$ and $\Delta u_B{}''(\phi) \leq 0$. The policymaker needs to choose which Policy ϕ to adopt, i.e., to choose the optimal value of ϕ. With a utilitarian social welfare function, the policymaker will choose ϕ to maximize: $\Delta u_A(\phi) + \Delta u_B(\phi)$.

2.2 Willingness to Pay (WTP)

Individual A is willing to pay m_A for a policy that increases her utility by Δu_A. And Individual B is willing to pay m_B for a policy that increases his utility by Δu_B. We thus have: $m_A(\Delta u_A, \omega_A)$ and $m_B(\Delta u_B, \omega_B)$. Since $\Delta u_A \geq 0$ and $\Delta u_B \geq 0$, we have $m_A(\Delta u_A, \omega_A) \geq 0$ and $m_B(\Delta u_B, \omega_B) \geq 0$.[30] Let $v(\omega)$ represent universal utility from wealth, namely, from purchasing a numeraire good. We assume that $v(\omega)$ is defined on $\Re^+$ and that $v(0) = 0$. We also assume decreasing marginal utility from wealth, i.e., $v'(\omega) > 0$ and $v''(\omega) < 0$. An individual $i \in \{A, B\}$ with wealth ω_i would divide this wealth between the policy change and the numeraire good. In particular, this individual's WTP, m_i, for a policy change that gives the individual Δu_i, is implicitly defined by:

$$(1) \quad v(\omega_i) - v(\omega_i - m_i) = \Delta u_i$$

as long as Equation (1) has a solution.[31] We have two ranges:

1) If $v(\omega_i) < \Delta u_i$, then $m_i = \omega_i$. In this range, Equation (1) does not have a solution; rather WTP for the policy change is determined by the individual rationality (IR) constraint: $m_i \leq \omega_i$ (which states that the individual would never be able to pay more than ω_i). If the utility from the policy change exceeds the utility obtained when the individual's wealth is spent entirely on the numeraire good, the individual would be willing to pay her entire wealth for the policy change.[32]

2) If $v(\omega_i) \geq \Delta u_i$, then Equation (1) has a solution and WTP for the policy change, m_i, is implicitly defined by Equation (1). If the utility from the policy change is smaller than the utility obtained when the individual's wealth is spent entirely on the numeraire good, then we get an interior solution that is implicitly defined by Equation (1).

WTP is an increasing function of the utility change and of wealth: $\frac{\partial m_i}{\partial \Delta u_i} \geq 0$ and $\frac{\partial m_i}{\partial \omega_i} \geq 0$. (From Equation (1), we know that $\frac{dm_i}{d\Delta u_i} = \frac{1}{v\prime(\omega_i - m_i)} > 0$ and $\frac{dm_i}{d\omega_i} = \frac{v\prime(\omega_i - m_i) - v\prime(\omega_i)}{v\prime(\omega_i - m_i)} > 0$.)

There are two types of WTP measures: individualized WTP and uniform, average WTP. In theory, we need to consider the individualized WTP – the WTP of the individuals who are affected by the policy. This means that when considering Policy A, we assess the WTP of the rich Individual A, $m_A(\Delta u_A, \omega_A)$; and when considering Policy B, we assess the WTP of the poor Individual B,

[30] If a policy increases individual *i*'s wealth and the individual can borrow against his future wealth, then individual *i*'s WTP for the policy will reflect the increased wealth. Indeed, ω_i properly understood should reflect the individual's future wealth. Compare: Kronman, *supra* note 1, at 240-241.

[31] This formulation assumes that the individual's overall utility is equal to the sum of her utility from the numeraire good, $v(\omega_i)$, and her utility from the policy, Δu_i. In a more general formulation, we would have utility $U(\omega_i, P)$ that is a general function of wealth and policy. And the WTP, m_i, for a policy change, from P^0 to P^1, would be implicitly defined by: $U(\omega_i, P^0) = U(\omega_i - m_i, P^1)$. Compare: Lewis A. Kornhauser, "On Justifying Cost-Benefit Analysis," 29 J. Legal Stud. 1037, 1040 (2000).

[32] More realistically, the maximum amount that an individual would be willing to pay for the policy change is not the individual's entire wealth, but rather it would be her entire wealth minus a certain amount that is needed to cover basic expenses (like housing, food, clothing, etc').

$m_B(\Delta u_B, \omega_B)$. This individualized measure, or at least WTP measures that are disaggregated by income groups, are sometimes used in practice.

In most cases, however, policymakers use a uniform, average (or median) WTP, aggregated across the entire population, even when the policy affects only a subset of the population. Most prominently, the value of a statistical life (VSL) that is routinely used in the cost-benefit analysis of regulations that affect mortality risk is a uniform, population-wide figure. The VSL is not calculated separately for each affected individual or even for each income group. Rather, a single VSL figure is used, averaging across the WTP of the rich and the poor (i.e., the WTP for a reduction in mortality risk). The VSL case represents a universal benefit – reduction in mortality risk – that, at least at an abstract level, provides a similar increase in utility for both the rich and the poor: $\Delta u_A = \Delta u_B \equiv \Delta u$. If Individual A's WTP for this benefit is $m_A(\Delta u, \omega_A)$ and Individual B's WTP for the same benefit is $m_B(\Delta u, \omega_B)$, then the average WTP is: $\overline{m}(\Delta u) = \frac{1}{2}\left(m_A(\Delta u, \omega_A) + m_B(\Delta u, \omega_B)\right)$. Such a uniform, average WTP will be used when assessing Policy A that affects only Individual A or Policy B that affects only Individual B.

2.3 Policy Effects and Policy Choices

Policymakers face different types of policy choices. In Case 1, the policymaker needs to decide whether or not to adopt a specific policy. The considered policy can be a policy that affects only the rich Individual A (Policy A) or a policy that affects only the poor Individual B (Policy B). If policymaking is based on WTP, then Policy A will be adopted *iff* m_A exceeds the cost of the policy, and Policy B will be adopted *iff* m_B exceeds the cost of the policy. In Case 2, the policymaker needs to choose between different policies with identical costs (that are funded from the same source, e.g., general tax revenues). Here, I assume that the policymaker must choose one policy from a list of proposed policies. The list may include policies that affect only the rich Individual A and policies that affect only the poor Individual B.[33] If policymaking is based on WTP, the policymaker will choose the policy with the highest WTP. Specifically, if the policymaker is choosing between Policy A and Policy B, then Policy A will be chosen if $m_A > m_B$ and Policy B will be chosen if $m_B > m_A$. (If $m_A = m_B$, then either policy can be chosen.)

2.4 The Informational Content of WTP

In a welfarist framework, the policymaker cares about preferences and utilities. An individual's WTP is relevant only to the extent that it contains information about that individual's utility. A perfectly informative measure will prefer Policy A over Policy B only when Policy A creates more utility for Individual A than Policy B creates for Individual B. An imperfectly informative measure might prefer Policy A even when it creates less utility. If WTP were perfectly informative, then $m_A > m_B$ only when $\Delta u_A > \Delta u_B$. With an imperfectly informative WTP, we get a distortion: $m_A > m_B$ even though $\Delta u_A < \Delta u_B$. (The preceding analysis may seem utilitarian, not just

[33] This type of policy choice can be motivated by a budget constraint that allows the policymaker to choose only one policy, especially when we are considering policies that are not paid-for by the individuals who benefit from the policy. Or by a notion of regulatory burden that limits the number of policies that can be adopted.

welfarist; but it isn't. Informational content is only one aspect of an overall welfare assessment. For example, with an egalitarian social welfare function, Policy B may be preferred, even if $\Delta u_A > \Delta u_B$ as perfectly indicated by $m_A > m_B$.)

We formalize the notion of informational content by defining and measuring the distortion caused by an imperfectly informative measure like WTP. Policy A *should* be chosen iff $\Delta u_A > \Delta u_B$, i.e., iff $\delta > 1$ (recall that $\Delta u_A = \delta \Delta u_B$). When WTP is imperfectly informative, there will be a threshold, $\hat{\delta} < 1$, such that for $\delta \in (\hat{\delta}, 1)$, WTP-based policymaking leads us astray: $m_A > m_B$ even though $\Delta u_A < \Delta u_B$. The threshold, $\hat{\delta}$, is implicitly defined by: $m_A(\Delta u_A = \hat{\delta}\Delta u_B, \omega_A) = m_B(\Delta u_B, \omega_B)$. Individual A is willing to pay more for Policy A than Individual B is willing to pay for Policy B as long as $\Delta u_A > \hat{\delta}\Delta u_B$.[34] The distortion, *D*, caused by the wealth disparity – the maximal difference between Δu_B and Δu_A (relative to Δu_B), for which Policy A will still be wrongly preferred over Policy B is:

$$D = \frac{\Delta u_B - \Delta u_A}{\Delta u_B} = \frac{\Delta u_B - \hat{\delta}\Delta u_B}{\Delta u_B} = 1 - \hat{\delta}$$

Another measure of the distortion is $1/\hat{\delta}$: WTP-based policymaking will prescribe policies that produce utility $1/\hat{\delta}$ times smaller than the alternative. With both measures, $1 - \hat{\delta}$ and $1/\hat{\delta}$, the distortion increases when $\hat{\delta}$ decreases. And, since a smaller distortion means larger informational content, we can measure the informational content of the WTP measure by $\hat{\delta}$.

This methodology for conceptualizing and quantifying the distortion caused by the WTP measure applies in both Case 1 and Case 2. In Case 1, each policy is evaluated independently, comparing the benefit from the policy as measured by WTP to the cost of the policy. With a perfectly informative measure, if Policy B that creates utility Δu_B is rejected, then Policy A that creates a smaller utility Δu_A will also be rejected. With an imperfectly informative WTP, Policy A might be adopted when Policy B is rejected, even when the utility created by Policy A is $1/\hat{\delta}$ times smaller. In Case 2, the policymaker chooses between Policy A and Policy B. With a perfectly informative measure, Policy A will never be chosen if it creates less utility than Policy B. With an imperfectly informative WTP, Policy A might be chosen over Policy B, even when the utility created by Policy A is $1/\hat{\delta}$ times smaller.

There is an alternative methodology for evaluating the distortion caused by an imperfectly informative WTP: Consider two policies, Policy A and Policy B, that create the same benefit, Δu, and measure the difference between Individual A's WTP for Policy A, $m_A(\Delta u, \omega_A)$, and Individual B's WTP for Policy B, $m_B(\Delta u, \omega_B)$. This difference represents a range of policies with a cost $c \in [m_B(\Delta u, \omega_B), m_A(\Delta u, \omega_A))$ that will be adopted when they benefit the rich, but not when the same benefit is enjoyed by the poor. In percentage terms, the distortion is: $D = \frac{m_A(\Delta u, \omega_A) - m_B(\Delta u, \omega_B)}{m_B(\Delta u, \omega_B)}$. A variation on this alternative measure proves especially useful, when we consider the informational

[34] Recall that Individual A's WTP for Policy A is $m_A(\Delta u_A, \omega_A)$ and Individual B's WTP for Policy B is $m_B(\Delta u_B, \omega_B)$. Since $\omega_A > \omega_B$ and $\frac{\partial m_i}{\partial \omega_i} \geq 0$, for $\Delta u_A = \Delta u_B$, or $\delta = 1$, we get $m_A(\delta \Delta u_B, \omega_A) > m_B(\Delta u_B, \omega_B)$. Since $\frac{\partial m_i}{\partial \Delta u_i} \geq 0$, there exists a threshold value $\hat{\delta} < 1$, such that $m_A(\hat{\delta}\Delta u_B, \omega_A) = m_B(\Delta u_B, \omega_B)$.

content of a uniform, average WTP measure in scenarios where policies are paid-for by the affected individuals. A uniform, average WTP is used when the policy creates a common benefit, Δu, like a reduction in mortality rate. The distortion occurs when a poor Individual B who is willing to pay $m_B(\Delta u, \omega_B)$ for the benefit is forced to pay the higher, average WTP, $\bar{m} = \frac{1}{2}[m_A(\Delta u, \omega_A) + m_B(\Delta u, \omega_B)]$; this distortion is measured by: $D_B = \frac{\bar{m} - m_B}{m_B} = \frac{1}{2}\left(\frac{m_A}{m_B} - 1\right)$. A parallel distortion occurs when a rich Individual A who is willing to pay $m_A(\Delta u, \omega_A)$ for the benefit is denied this benefit, because the policymaker is using the lower, average WTP, $\bar{m}$; this distortion is measured by: $D_A = \frac{m_A - \bar{m}}{m_A} = \frac{1}{2}\left(1 - \frac{m_B}{m_A}\right)$.

3. Policies that are Not Paid-For by the Affected Individuals

I begin by considering policies that are not paid-for by the affected individuals, e.g., policies that are funded by general tax revenues. Section 3.1 focuses on individualized WTP and studies the relationship between informational content and wealth disparity. Section 3.2 focuses on uniform, average WTP and shows that this measure reduces the distortion caused by wealth disparity.

3.1 Individualized WTP

3.1.1 Wealth Disparity and Informational Content

If Individual A and Individual B have the same wealth, $\omega_A = \omega_B$, then A's WTP will exceed B's WTP, i.e., $m_A \geq m_B$ iff $\Delta u_A \geq \Delta u_B$.[35] The policy that creates the largest increase in utility will be chosen. However, the greater the wealth disparity, the more likely it is that $m_A > m_B$ even though $\Delta u_A < \Delta u_B$. The relationship between the degree of wealth disparity (γ) and the informational content of WTP ($\hat{\delta}$) is summarized in the following proposition:

Proposition 1: The threshold, $\hat{\delta}$, is decreasing in the degree of wealth disparity, γ. Therefore, a less equal wealth distribution reduces the informational content of WTP.

Proof: The threshold $\hat{\delta}$ is implicitly defined by the equation: $m_A\left(\Delta u_A = \hat{\delta}\Delta u_B, \omega_A = \gamma\omega_B\right) = m_B(\Delta u_B, \omega_B)$. Taking the derivative of this equation w.r.t. γ, we obtain: $\frac{dm_A}{d(\hat{\delta}\Delta u_B)} \cdot \Delta u_B \cdot \frac{d\hat{\delta}}{d\gamma} + \frac{dm_A}{d(\gamma\omega_B)} \cdot \omega_B = 0$. Or: $\frac{d\hat{\delta}}{d\gamma} = -\frac{dm_A}{d(\gamma\omega_B)} \cdot \frac{\omega_B}{\Delta u_B} \cdot \left(\frac{dm_A}{d(\hat{\delta}\Delta u_B)}\right)^{-1}$. Eq. (1) implies: $\frac{dm_i}{d\Delta u_i} = \frac{1}{v\prime(\omega_i - m_i)} > 0$ and $\frac{dm_i}{d\omega_i} = 1 - \frac{v'(\omega_i)}{v'(\omega_i - m_i)} > 0$. Therefore, $\frac{dm_A}{d(\gamma\omega_B)} > 0$ and $\frac{dm_A}{d(\hat{\delta}\Delta u_B)} > 0$. And we have: $\frac{d\hat{\delta}}{d\gamma} < 0$. QED

[35] This result follows from: (i) Given $\omega_A = \omega_B$, under Equation (1) $\Delta u_A = \Delta u_B$ implies $m_A = m_B$, and (ii) Equation (1) implies $\frac{dm_i}{d\Delta u_i} \geq 0$.

To summarize: With individualized WTP measures, when wealth disparity is large, the effect of wealth on WTP dilutes the effect of preferences on WTP reducing the informational content of WTP.

3.1.2 Example

To get a better handle on the size of the distortion caused by WTP-based policymaking, we add some additional structure to the model. We measure the utilities from the policy change in relation to the utility that Individual B would have received had she spent her entire wealth on the numeraire good. Specifically, we have $\Delta u_B = \alpha v(\omega_B)$, where $\alpha \geq 0$. A higher α represents policies with bigger impact, and a lower α represents policies with smaller impact. For Individual A, using $\Delta u_A = \delta \Delta u_B$, we have $\Delta u_A = \delta \alpha v(\omega_B)$. Finally, we consider a specific functional form: $v(\omega) = \beta \ln \omega$, which has been commonly used in the literature.[36] Based on subjective well-being (or happiness) data from Layard et al (2008), we set $\beta = \frac{1}{2}$, so that $v(\omega) = \frac{1}{2} \ln \omega$.[37]

With this additional structure, we can derive a closed-form expression for the threshold $\hat{\delta}$ (see Appendix). Specifically –

$$(2) \quad \hat{\delta} = \begin{cases} \frac{1}{2\alpha} ln\left(\frac{\gamma}{(\gamma-1)+e^{-2\alpha}}\right) & , \quad \alpha < 1 + \frac{1}{2}\ln \omega_B \\ \frac{1}{2\alpha} ln\left(\frac{\gamma \omega_B}{(\gamma-1)\omega_B+e^{-2}}\right) & , \quad \alpha \geq 1 + \frac{1}{2}\ln \omega_B \end{cases}$$

In Figure 1 (below), we graph the threshold $\hat{\delta}$ as a function of α (the solid black line).

[36] See Bernoulli, D., 1738. Specimen theoriae novae de mensura sortis. Commentarii Academiae Scientiarum Imperialis Petropolitanae (for 1730 and 1731) 5, 175–192; Bernoulli, D., 1954. Exposition of a New Theory on the Measurement of Risk. Econometrica 22, 23–36; Dalton, H., 1920. The measurement of the inequality of incomes. Economic Journal 30 (119), 348–361; Layard, R, S. Nickell, and G. Mayraz. 2008. The Marginal Utility of Income. J. Public Econ. 92: 1846-1857. The Appendix analyzes an alternative functional form and explores the sensitivity of the quantified distortion to the chosen function.

[37] Setting $\beta = \frac{1}{2}$ is also consistent with Viscusi and Masterman's estimates of the income elasticity of the VSL (Value of a Statistical Life). See W. Kip Viscusi and Clayton J. Masterman, "Income Elasticities and Global Values of a Statistical Life," 8 J. of Benefit-Cost Analysis 226 (2017).

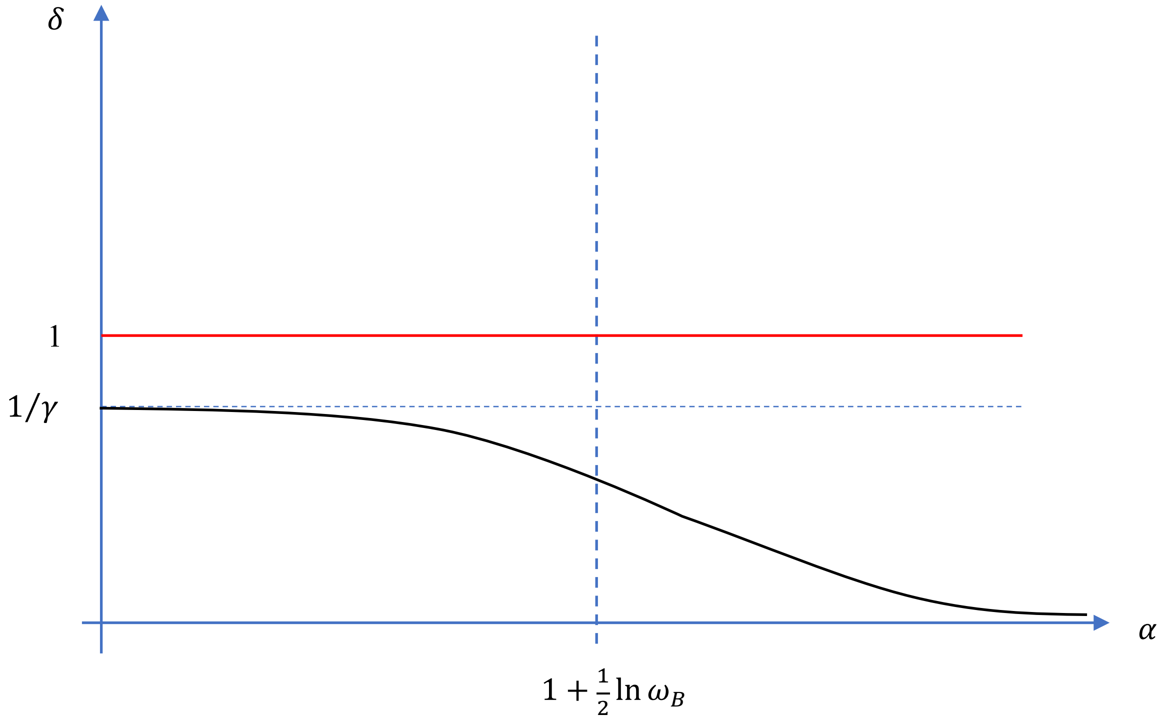

Figure 1: The threshold $\hat{\delta}$ as a function of α

Policy A *should* be chosen (or preferred over Policy B) in the area above the red line, i.e., when $\delta > 1$. With WTP-based policymaking, Policy A *will* be chosen (or preferred over Policy B) in the area above the black line, i.e., when $\delta > \hat{\delta}$. The distortion zone is the area between the black and red lines. The size of this zone increases with the magnitude of the wealth disparity (γ). The distortion, measured by the distance between the black and red lines, is increasing in the impact of the policy (α). With a bigger policy change, the effect of wealth is larger. And the distortion caused by wealth disparity is correspondingly larger.

Some simple comparative statics can help to ascertain the size of the distortion that is caused by wealth disparity. Starting with the magnitude of the wealth disparity, γ: What is a plausible value for this parameter? Since the simple example includes only two representative individuals, A and B, with different levels of wealth, we let A represent an individual at the 60th percentile of wealth and we let B represent an individual at the 40th percentile of wealth. In the US, this means that A's wealth is ~\$170,000 or, measured relative to the median wealth of ~\$100,000, $\omega_A = 1.7$; and B's wealth is ~\$50,000 or, measured relative to the median wealth of ~\$100,000, $\omega_B = 0.5$.[38] This implies: $\gamma = \frac{\omega_A}{\omega_B} = 3.4$. Using these parameter values, the following table presents the

[38] The wealth figures are taken from PK, United States Net Worth Brackets, Percentiles, and Top One Percent (https://dqydj.com/net-worth-brackets-wealth-brackets-one-percent/) (Based on the 2016 Surveys of Consumer Finances from the Federal Reserve).

informational content ($\hat{\delta}$) and the distortion measure ($1/\hat{\delta}$), for policy changes with different impact levels, i.e., with different α values.

α	$\hat{\delta}$	$1/\hat{\delta}$
0.01	0.29	3.44
0.1	0.27	3.70
0.25	0.25	4.00
0.5	0.20	5.00
1	0.12	8.33
1.5	0.08	12.5
2	0.06	16.67

Table 2(a): Distortion Caused by Wealth Disparity – Smaller Disparity ($\gamma = 3.4$, $\omega_B = 0.5$)

We see that for a smaller-impact policy, e.g., $\alpha = 0.01$, a WTP-based assessment will prescribe (or prefer) Policy A even if it produces utility that is 3.44 smaller than Policy B; and for a larger-impact policy, e.g., $\alpha = 2$, a WTP-based assessment will prescribe (or prefer) Policy A even if it produces utility that is 16.67 times smaller than Policy B.

What happens if we increase the wealth disparity? Let A represent an individual at the 70th percentile of wealth and we let B represent an individual at the 30th percentile of wealth. In the US, this means that A's wealth is $\sim\$280{,}000$ or, measured relative to the median wealth of ~\$100,000, $\omega_A = 2.8$; and B's wealth is $\sim\$19{,}000$ or, measured relative to the median wealth of ~\$100,000, $\omega_B = 0.19$. This implies: $\gamma = \frac{\omega_A}{\omega_B} = 14.7$. Using these parameter values, the following table presents the informational content ($\hat{\delta}$) and the distortion measure ($1/\hat{\delta}$), for policy changes with different impact levels, i.e., with different α values.

α	$\hat{\delta}$	$1/\hat{\delta}$
0.01	0.07	14.29
0.1	0.06	16.67
0.25	0.04	25.00
0.5	0.02	50.00
1	0.01	100
1.5	0.007	142.86
2	0.005	200.00

Table 2(b): Distortion Caused by Wealth Disparity – Larger Disparity ($\gamma = 14.7$, $\omega_B = 0.19$)

We see that for a smaller-impact policy, e.g., $\alpha = 0.01$, a WTP-based assessment will prescribe (or prefer) Policy A even if it produces utility that is 14.29 smaller than Policy B; and for a larger-impact policy, e.g., $\alpha = 2$, a WTP-based assessment will prescribe (or prefer) Policy A even if it produces utility that is 200 times smaller than Policy B. As the wealth disparity increases – moving from Table 2(a) ($\gamma = 3.4$) to Table 2(b) ($\gamma = 14.7$) – the magnitude of the distortion increases. The increase in the distortion is roughly proportional to the increase in wealth disparity for smaller-impact policies, and more than proportional for larger-impact policies.

3.1.3 A Wealth-Adjusted WTP

I used the proposed analytical framework to quantify the distortion caused by wealth disparity. The same analytical framework can be used to correct for the wealth disparity, namely, to derive a wealth-adjusted WTP. Consider the poor Individual B, with wealth ω_B, who is willing to pay m_B for a policy that creates a benefit Δu_B. How much would an individual with median wealth pay for the same benefit? The WTP of the median individual can be thought of as a wealth-adjusted WTP.

In our simple, two-person example, the median wealth is: $\overline{\omega} = \frac{1}{2}(\omega_A + \omega_B) = \frac{1}{2}(1+\gamma)\omega_B$. Let $m(\overline{\omega})$ denote the median-individual's WTP for a benefit Δu_B. We can derive the ratio $\frac{m(\overline{\omega})}{m_B}$ for policies or benefits of different magnitude, represented by the parameter α (as in the preceding analysis). This ratio is derived in the Appendix. The individual WTP would have to be "scaled up" or multiplied by $\frac{m(\overline{\omega})}{m_B}$. Table 3 shows the required $\frac{m(\overline{\omega})}{m_B}$ for different α values. When the wealth disparity is smaller ($\gamma = 3.4, \omega_B = 0.5$), the poor individual's WTP needs to be multiplied by 2.2 – 2.64, depending on the impact or magnitude of the policy. When the wealth disparity is larger ($\gamma = 14.7, \omega_B = 0.19$), the poor individual's WTP needs to be multiplied by 7.85 – 24.64, depending on the impact or magnitude of the policy.

α	$\frac{m(\overline{\omega})}{m_B}$ $\gamma = 3.4$ and $\omega_B = 0.5$
$\alpha \leq 0.65$	2.20
0.7	2.27
0.8	2.41
0.9	2.52
1	2.61
$\alpha \geq 1.05$	2.64

(a)

α	$\frac{m(\overline{\omega})}{m_B}$ $\gamma = 14.7$ and $\omega_B = 0.19$
$\alpha \leq 0.17$	7.85
0.4	14.92
0.6	18.93
0.8	21.62
1.0	23.42
$\alpha \geq 1.20$	24.64

(b)

Table 3: Wealth Multipliers for WTP – (a) Smaller Wealth Disparity; (b) Larger Wealth Disparity

We could apply a similar conversion to the WTP of the rich, Individual A. By using these wealth-adjusted WTP figures, policymakers can avoid the distortion caused by wealth disparity. To calculate the wealth-adjusted WTP, policymakers need to know the wealth distribution and the impact or magnitude of the policies under consideration.

3.2 Uniform, Average WTP

Shifting from individualized WTP to average, uniform WTP avoids distortions in policymaking, when it is used to assess universal benefits, like reduction in mortality risk, that provide a similar increase in utility for both the rich and the poor: $\Delta u_A = \Delta u_B = \Delta u$. Consider Policy A that reduces mortality risk for Individual A by a multiple $a > 0$, such that $\Delta u_A = a\Delta u$, and Policy B that reduces mortality risk for Individual B by a multiple $b > 0$, such that $\Delta u_B = b\Delta u$. With a uniform WTP, $\overline{m}$, for a reduction in mortality risk, Δu, WTP is informative: Policy B will be preferred *iff* $b > a$.

3.3 Summary

When the considered policies are *not* paid for by the affected individuals, then the informational content of individualized WTP is decreasing in the degree of wealth disparity. The individualized WTP distorts policymaking in a particular direction – benefiting the rich at the expense of the poor. A uniform, average WTP has more informational content, at least when the considered policies create universal benefits, like reduction in mortality risk. Moreover, in the important case of universal benefits, a uniform, average WTP can support progressive redistribution, independent of informational content. Consider a policy that saves many (statistical) lives of poor individuals, but costs billions to implement. Using the poor individuals' WTP, the policymaker might conclude that the benefit does not justify the cost and reject the policy. Using the higher, average WTP, the same policy would be adopted. (Using the average WTP and adopting the policy is especially good for the poor, if the implementation costs are paid for by general taxes and the poor pay less taxes.) Now consider a policy that saves many (statistical) lives of rich individuals. Using the high WTP of the rich, the policy would be adopted, despite high implementation costs. The same policy might be rejected if we use the lower, average WTP.

4. Policies that are Paid-for by the Affected Individuals

I now consider policies that are paid-for by the affected individuals. Section 4.1 focuses on individualized WTP and shows that this measure has high informational content. Section 4.2 focuses on uniform, average WTP, shows that this measure has lower informational content and studies the relationship between informational content and wealth disparity.

4.1 Individualized WTP

When a policy creates a benefit and this benefit is paid for by the affected individual, the individualized WTP perfectly balances benefit and cost. The poor are willing to pay less than the rich for a policy that would create the same (or greater) utility, because the poor have other, high-utility uses for the little money that they have, e.g., paying rent and buying food. The rich, on the other hand, have more money and lower utility uses for the marginal dollars (think of a billionaire buying yet another yacht.) WTP thus balances the utility created by the policy (- the benefit side) against the utility from alternative uses (- the cost side). Since both benefits and costs are important, WTP is normatively appealing. In some sense, the poor really want the policy less than the rich. WTP is normatively appealing, when the cost of implementing the policies is borne by the individuals who are affected by the policy. In this case, it makes sense to adopt a policy that affects the rich and reject an equal-benefit policy that affects the poor. The individualized WTP is informative.

4.2 Uniform, Average WTP

4.2.1 Wealth Disparity and Informational Content

When policies are paid-for by the affected individuals, a uniform, average WTP, like the VSL, has low informational content. It distorts policy choice, and this distortion increases with the degree of wealth disparity. I quantify these distortions by showing how greater wealth disparity increases the range of welfare-reducing policies that would be adopted if average WTP is used.

I consider two policies, Policy A that creates a benefit Δu, i.e., reduction in mortality risk, for the rich Individual A, and Policy B that creates an equivalent benefit Δu for the poor Individual B. There are two types of distortion (see Section 2): (1) Welfare-reducing adoption of Policy B – a distortion measured by $D_B = \frac{\bar{m}-m_B}{m_B} = \frac{1}{2}\left(\frac{m_A}{m_B} - 1\right)$; and (2) Welfare-reducing failure to adopt Policy A – a distortion measured by $D_A = \frac{m_A-\bar{m}}{m_A} = \frac{1}{2}\left(1 - \frac{m_B}{m_A}\right)$. If Individual A and Individual B have the same wealth, $\omega_A = \omega_B$, then $\bar{m} = m_A = m_B$ and thus there are no distortions ($D_B = D_A = 0$). As the wealth disparity increases, the distortion increases. The relationship between the degree of wealth disparity (γ) and the distortions (D_B and D_A) is summarized in the following proposition:

Proposition 2: The distortions D_B and D_A are increasing in the degree of wealth disparity, γ. Therefore, a less equal wealth distribution reduces the informational content of WTP.

Proof: Individual A's WTP, $m_A(\Delta u, \gamma\omega_B)$, is increasing in γ: $\frac{dm_A}{d\gamma} > 0$ (since $\frac{dm_i}{d\omega_i} > 0$). Therefore, $\frac{dD_B}{d\gamma} = \frac{1}{2m_B} \cdot \frac{dm_A}{d\gamma} > 0$, and $\frac{dD_A}{d\gamma} = \frac{m_B}{2(m_A)^2} \cdot \frac{dm_A}{d\gamma} > 0$. QED

4.2.2 Example

I use a variation on the example from Section 3.1.2. For concreteness, consider the benefit from a reduction in mortality risk. Let $\Delta u = \alpha v(\omega_B)$, with $\alpha \geq 0$, denote this benefit, which is assumed to be independent of wealth. A higher α represents policies with bigger impact; here, a larger reduction in mortality risk. Individualized WTP measures, m_A for Individual A and m_B for Individual B, are derived as before. Here, the utility change, Δu, is identical for both parties, but still the wealth disparity results in different individualized WTP values.

(1) Welfare-Reducing Adoption of Policy B

Policy B *should* be adopted when the cost, which is also the price that Individual B will pay (e.g., for the safer product), is smaller than m_B. If a uniform, average WTP is used, then Policy B *will* be adopted when the cost is smaller than $\bar{m}$ (which is larger than m_B). The magnitude, in percentage terms, of the distortion is: $D_B = \frac{\bar{m}-m_B}{m_B}$. Plugging in the expressions for m_B and $\bar{m}$, which are derived in the Appendix, we obtain:

$$D_B = \frac{\bar{m}-m_B}{m_B} = \begin{cases} \dfrac{\frac{1}{2}(\gamma-1)\omega_B}{\omega_B - 0.135} & , \quad \alpha \geq \hat{\alpha}_A \\ \dfrac{\frac{1}{2}[\gamma\omega_B(1-e^{-2\alpha}) - (\omega_B - 0.135)]}{\omega_B - 0.135} & , \quad \hat{\alpha}_B \leq \alpha < \hat{\alpha}_A \\ \frac{1}{2}(\gamma-1) & , \quad \alpha < \hat{\alpha}_B \end{cases}$$

where $\hat{\alpha}_B \equiv 1 + \frac{1}{2}\ln\omega_B$ and $\hat{\alpha}_A \equiv 1 + \frac{1}{2}\ln\gamma\omega_B$.

Figure 2 plots the distortion, D_B, as a function of the magnitude of the benefit, α.

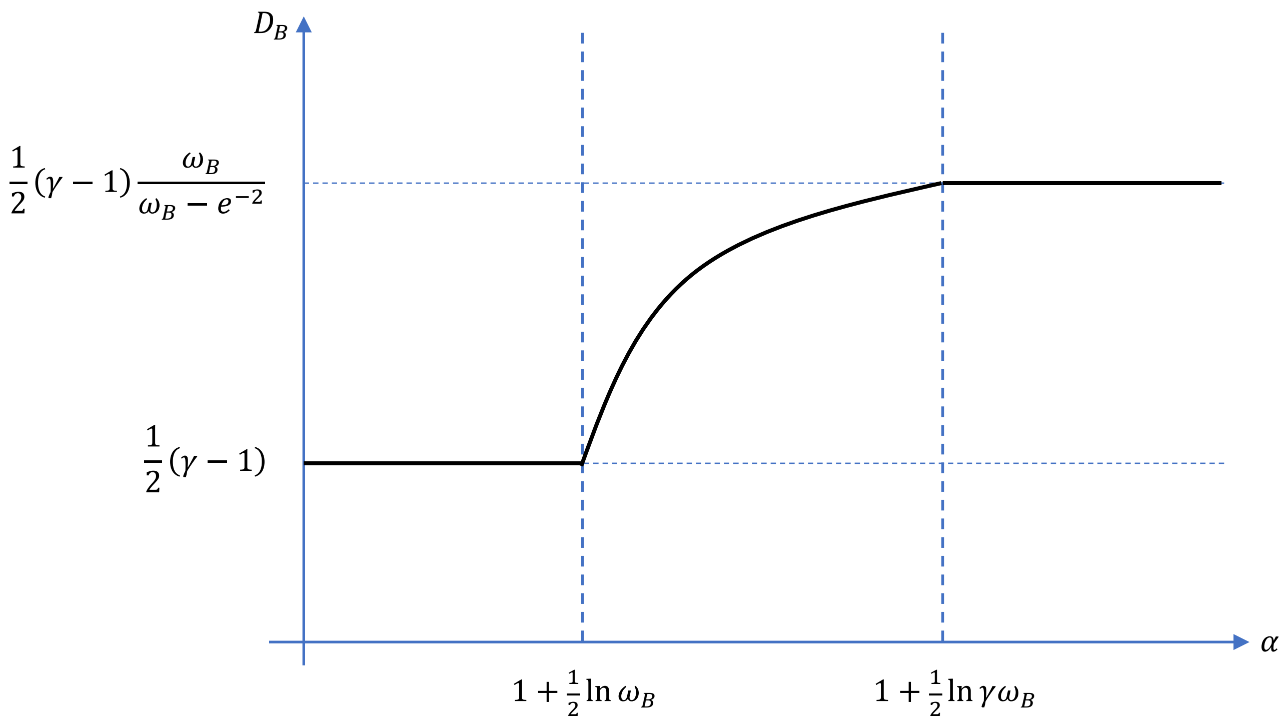

Figure 2: Type 1 distortion, D_B, as a function of α

We see that the distortion is (weakly) increasing in the magnitude of the benefit, namely, in the size of the reduction in mortality risk, α. More important, we see that the distortion is increasing in the degree of wealth disparity, γ.

Some simple comparative statics can help to ascertain the size of the distortion that is caused by wealth disparity. As before, let $\omega_B = 0.5$ and $\gamma = \frac{\omega_A}{\omega_B} = 3.4$. Let's consider policy changes with different impact levels, i.e., with different α values. For small-impact policies, with $\alpha < 1 + \frac{1}{2}\ln \omega_B = 0.65$, the distortion is $D = \frac{1}{2}(\gamma - 1) = 1.2$, namely, Individual B would be forced to pay 20% more than he is willing to pay for this benefit. For large-impact policies, with $\alpha > 1 + \frac{1}{2}\ln \gamma\omega_B = 1.27$, the distortion is $D = \frac{1}{2}(\gamma - 1)\frac{\omega_B}{\omega_B - e^{-2}} = 1.64$, namely, Individual B would be forced to pay 64% more than he is willing to pay for this benefit. For intermediate-impact policies, the distortion would be between 20% and 64%.

What happens if we increase the wealth disparity and set $\omega_B = 0.19$ and $\gamma = \frac{\omega_A}{\omega_B} = 14.7$? For small-impact policies, with $\alpha < 1 + \frac{1}{2}\ln \omega_B = 0.17$, the distortion is $D = \frac{1}{2}(\gamma - 1) = 6.85$, namely, Individual B would be forced to pay 6.85 times as much as the benefit is actually worth to him, a distortion of 585%. For large-impact policies, with $\alpha > 1 + \frac{1}{2}\ln \gamma\omega_B = 1.51$, the distortion is $D = \frac{1}{2}(\gamma - 1)\frac{\omega_B}{\omega_B - e^{-2}} = 9.38$, namely, Individual B would be forced to pay 9.38 times as much as the benefit is actually worth to him, a distortion of 838%. For intermediate-impact policies, the distortion would be between 585% and 838%.

(2) Welfare-Reducing Failure to Adopt Policy A

Policy A *should* be adopted when the cost, which is also the price that Individual A will pay (e.g., for the safer product), is smaller than m_A. If a uniform, average WTP is used, then Policy A *will* be adopted when the cost is smaller than $\bar{m}$ (which is smaller than m_A). The magnitude, in percentage terms, of the distortion is: $D_A = \frac{m_A - \bar{m}}{m_A}$. We can derive expressions for m_A and $\bar{m}$ and quantify the distortion caused by the failure to adopt Policy A, as we did with the distortion caused by welfare-reducing adoption of Policy B. The analysis is very similar and is, therefore, omitted.

4.3 Summary

When the considered policies are paid for by the affected individuals, then individualized WTP has high informational content. It is noteworthy, however, that this high informational content does not prevent individualized WTP from increasing wealth disparity. There will be a range of expensive policies that will be adopted when they benefit the rich, but not when they benefit the poor. These policies will further advantage the rich relative to the poor. A uniform, average WTP has lower informational content, and its informational content is decreasing with the degree of wealth disparity. The uniform, average WTP harms both the rich and the poor, resulting in ambiguous distributional implications.

5. Time and Rationality

5.1 Time

Thus far, we have considered a static, one-period framework. We now consider a dynamic extension with multiple time periods. Utilities, social welfare and wealth, change over time, as new policies are adopted. Let u_i^t denote the utility of individual i at time *t*. Let $W^t = W(u_1^t, u_2^t, \dots, u_N^t)$ denote social welfare at time *t*. And let $\overline{\omega}^t = (\omega_1^t, \omega_2^t, \dots, \omega_N^t)$ denote the vector of wealth values at time *t*. Given the initial wealth distribution, $\overline{\omega}^0$, a WTP-based analysis, conducted at t=1, leads to the adoption of policy P^1. Policy P^1 then changes the wealth distribution to $\overline{\omega}^1$. With $\overline{\omega}^1$, a WTP-based analysis, conducted at t=2, leads to the adoption of policy P^2. Policy P^2 then changes the wealth distribution to $\overline{\omega}^2$. Etc.

A standard critique of WTP-based policymaking is that the chosen policy depends on the initial distribution of wealth. The dynamic extension strengthens this critique. The initial wealth distribution affects not only the current policy choice, but also many future policy choices. If the initial wealth distribution, $\overline{\omega}^0$, is even slightly unequal, this inequality might result in a policy P^1 that slightly increases this inequality. The more skewed wealth distribution, $\overline{\omega}^1$, might result in a policy P^2 that further increases inequality. And so on. Through this dynamic, inequality can increase over time.

Related, and perhaps even more interesting, is the effect of such policy dynamics on the initial WTP-based analysis and thus on the initial policy decision. Since the initial policy will affect, through the evolving wealth distribution, many future policies, the stakes are higher and thus WTP for the initial policy will be higher. Indeed, individuals would borrow against future wealth to increase WTP and secure their favored (initial) policy. Moreover, the individual's t=0 wealth, properly understood, incorporates – in a Net Present Value (NPV) sense – anticipated future changes in her wealth, at least to the extent that the individual can borrow against her future wealth. Thus, even if two individuals start off with identical wealth, the possibility that successive policy choices could benefit one individual more than the other implies that wealth disparity might already distort the initial, t=0 policy choice. The dynamic extension suggests additional concerns about WTP-based policymaking.

5.2 Rationality

The preceding discussion assumed that individuals are rational, in the sense that they anticipate the consequences of the policy change and adjust their WTP accordingly. But some individuals might not be able to anticipate all possible consequences of a policy change and how it affects their utility going forward. And even if they anticipate these consequences, imperfectly rational individuals might struggle to translate them into a WTP. It is not surprising that such deviations from perfect rationality can skew WTP-based policymaking, and any other type of policymaking. Still, there are several specific implications of imperfect rationality that should be noted.

Start with the static framework and consider the standard WTP question: "how much are you willing to pay for a Policy X?" This question would elicit a response that is sensitive to the immediate effect of the policy change, e.g., "I think that rearview cameras should be added to cars and I would be willing to pay an extra $X for a car with a rearview camera." But what about the distributional effects of the policy? For example, a rule that requires rearview cameras in all cars would increase the price of cars, thus harming poorer individuals disproportionately. Either all car buyers would pay more, but the same price increase would have a larger effect on the poor. Or, the policy may have no effect at all on the rich, if the high-end cars that the rich buy were already equipped with rearview cameras (voluntarily, before the policy change mandated this feature). While a rational individual would consider these broader, distributional effects when stating a WTP for the policy, an imperfectly rational individual might not.

The rationality assumption becomes even more unrealistic in the dynamic extension. Now, beyond the immediate effect of the policy, and its distributional consequences, the individual would need to anticipate a progression of policy changes that would be triggered by the initial policy choice, and the greater distributional consequences of this progression. A perfectly rational individual would state a WTP that incorporates these long-term effects. An imperfectly rational, or myopic, individual will not.

It is not hard to imagine how imperfect rationality could distort WTP-based policymaking. For example, consider a choice between policy P^{1A} and policy P^{1B}. If P^{1A} is adopted, then P^{2A}, P^{3A}, etc. will follow; this will result in a direct increase in Individual A's utility and will also enrich Individual A. If P^{1B} is adopted, then P^{2B}, P^{3B}, etc' will follow; this will result in a direct increase in Individual B's utility and will also enrich Individual B. Now assume that the parties are differentially rational, such that Individual A fully anticipates the dynamic effects of choosing policy P^{1A}, whereas Individual B is myopic and considers only the immediate effects of policy P^{1B}. In this example, a WTP-based assessment might lead to the adoption of policy P^{1A}, even though P^{1B}, P^{2B}, P^{3B}, etc' would result in greater social welfare. The question of rationality raises additional concerns about WTP-based policymaking.[39]

6. Conclusion

This Article offered a welfarist reassessment of WTP-based policymaking. Most fundamentally, WTP-based assessments are justified, in a welafrist framework, to the extent that WTP provides information about individual preferences. The preceding analysis suggested potentially significant limits on the informational content of WTP, especially in a society with large wealth disparity. It also quantified the distortions that result, when policymakers use a WTP measure with low informational content. The informational content of WTP depends on the policymaking context and on the specific WTP measure that is applied. The results in this Article can help policymakers identify the best WTP measure for the policy choice that they face. In some scenarios, WTP is quite attractive. In others, the shortcomings of WTP-based policymaking are significant, and they should be considered by policymakers, before they use WTP to guide their policy choices.

---

[39] Indeed, the challenges that imperfect rationality poses for WTP-based policymaking go beyond those outlines here. *See* Sunstein, *supra* note **Error! Bookmark not defined.**, at 403, 411, 427-28.

Appendix

Section A1 of this Appendix develops the example from Section 3.1.2, using the functional form: $v(\omega) = \beta \ln \omega$. Using the same example, Section A2 derives the wealth-adjusted WTP for the Section 3.1.3 analysis. Section A3 develops another example using the alternative functional form: $v(\omega) = \sqrt{\omega}$. Section A4 develops the example from Section 4.2.2, using the functional form: $v(\omega) = \beta \ln \omega$.

A1. Section 3.1.2: $v(\omega) = \beta \ln \omega$; Individualized WTP

We follow Layard et al (2008) who use and justify the functional form, $v(\omega) = \beta \ln \omega$, as a first approximation. Based on happiness data summarized in Layard et al (2008), $\beta \approx 0.5$. So we have: $v(\omega) = \frac{1}{2} \ln \omega$.[40] Following Layard et al (2008), we interpret ω as an individual's wealth relative to the median wealth level in the population. For example, $\omega = 1$ represents an individual with a level of wealth that is equal to the median wealth level in the population, and $\omega = 0.5$ represents an individual with a level of wealth that is half of the median wealth level in the population. Correspondingly, the utility of the individual with the median wealth level in the population is normalized to zero, i.e., $v(1) = \frac{1}{2} \ln 1 = 0$.

We assume, without loss of generality, that A is richer than B and let $\omega_A = \gamma \omega_B$, with $\gamma > 1$. We measure the utilities from the policy change in relation to the difference between the utility of a person with average wealth ($\omega = 1$) and the utility of a person with subsistence wealth, which we set at 13.5% of the average wealth ($\omega = e^{-2} \approx 0.135$). Specifically, for individual B, we have $\Delta u_B = \alpha\left(\frac{1}{2} \ln 1 - \frac{1}{2} \ln 0.135\right) = \alpha$, where $\alpha \geq 0$. And, for Individual A, we have $\Delta u_A = \delta \Delta u_B$, where $\delta \geq 0$. Or, substituting $\Delta u_B = \alpha$, we have $\Delta u_A = \delta \alpha$.

Now consider the parties' WTP. We measure WTP, m, as a percent of the average wealth level, so that it can be easily compared to the wealth percentile (ω), as defined above. We assume that an individual will not pay an amount that would leave this individual with less than subsistence level wealth. (The subsistence level assumption avoids the technical difficulties of dealing with a logarithm of zero.) This implies: $\omega - m \geq 0.135$. And substituting into the equation: $\frac{1}{2} \ln \omega - \frac{1}{2} \ln(\omega - m) \leq \Delta u$, we get a corner solution of $m = \omega - 0.135$, whenever $\frac{1}{2} \ln \omega + 1 \leq \Delta u$.

For Individual B, we have

[40] The median β in Layard et al (2008), Table 3, is 0.57. More fundamentally, several assumptions – about the relationship between income and happiness and between happiness and utility – are necessary to justify the functional form: $v(\omega) = \beta \ln \omega$. [Cardinal utility, utility that is comparable across individuals, subjective well-being (or happiness) as a measure of utility, etc'.] See Layard et al. (2008) for a discussion these assumptions and their justification.

$$m_B = \begin{cases} \omega_B - 0.135 & , \quad \frac{1}{2}\ln \omega_B + 1 \le \Delta u_B \\ \omega_B(1 - e^{-2\Delta u_B}) & , \quad \frac{1}{2}\ln \omega_B + 1 > \Delta u_B \end{cases}$$

And, after substituting $\Delta u_B = \alpha$:

$$m_B = \begin{cases} \omega_B - 0.135 & , \quad \alpha \ge \hat{\alpha} \\ \omega_B(1 - e^{-2\alpha}) & , \quad \alpha < \hat{\alpha} \end{cases}$$

where $\hat{\alpha} \equiv 1 + \frac{1}{2}\ln \omega_B$. (Note that $\omega_B > 0.135$ implies $\hat{\alpha} > 0$.)

And for Individual A, we have

$$m_A = \begin{cases} \omega_A - 0.135 & , \quad \frac{1}{2}\ln \omega_A + 1 \le \Delta u_A \\ \omega_A(1 - e^{-2\Delta u_A}) & , \quad \frac{1}{2}\ln \omega_A + 1 > \Delta u_A \end{cases}$$

And, after substituting $\omega_A = \gamma\omega_B$ and $\Delta u_A = \delta\alpha$:

$$m_A = \begin{cases} \gamma\omega_B - 0.135 & , \quad \delta \ge \hat{\delta}_A \ or \ \delta\alpha \ge k(\omega_B, \gamma) \\ \gamma\omega_B\left(1 - e^{-2\delta\alpha}\right) & , \quad \delta < \hat{\delta}_A \ or \ \delta\alpha < k(\omega_B, \gamma) \end{cases}$$

where $k(\omega_B, \gamma) \equiv 1 + \frac{1}{2}\ln \gamma\omega_B$ and $\hat{\delta}_A = \frac{k(\omega_B,\gamma)}{\alpha}$ ($= \frac{1+\frac{1}{2}\ln \gamma\omega_B}{\alpha} = \frac{1}{2\alpha}(2 + \ln \gamma\omega_B) = \frac{1}{2\alpha}\ln \gamma\omega_B\, e^2$). (Note that $\omega_B > 0.135$ and $\gamma > 1$ imply $k(\omega_B, \gamma) > 0$ and $\hat{\delta}_A > 0$.)

Policy A *will* be chosen iff $m_A > m_B$.

There are four ranges:

(i) $\alpha < \hat{\alpha}$ and $\alpha\delta < k(\omega_B, \gamma)$: Policy A *will* be chosen iff $\gamma\left(1 - e^{-2\delta\alpha}\right) > 1 - e^{-2\alpha}$, or $\delta > \hat{\delta}_1 = \frac{1}{2\alpha} ln\left(\frac{\gamma}{(\gamma-1)+e^{-2\alpha}}\right)$. [Note: $\hat{\delta}_1(\alpha) < \hat{\delta}_A$]

(ii) $\alpha > \hat{\alpha}$ and $\alpha\delta < k(\omega_B, \gamma)$: Policy A *will* be chosen iff $\gamma\omega_B\left(1 - e^{-2\delta\alpha}\right) > \omega_B - 0.135$, or $\delta > \hat{\delta}_2 = \frac{1}{2\alpha} ln\left(\frac{\gamma\omega_B}{(\gamma-1)\omega_B+e^{-2}}\right)$. [Note: $\hat{\delta}_2(\alpha) < \hat{\delta}_A$.]

(iii) $\alpha < \hat{\alpha}$ and $\alpha\delta > k(\omega_B, \gamma)$: Policy A *will* be chosen iff $\gamma\omega_B - 0.135 > \omega_B(1 - e^{-2\alpha})$, which is always true (since $\omega_B - 0.135 > \omega_B(1 - e^{-2\alpha})$and $\gamma > 1$ by assumption).

(iv) $\alpha > \hat{\alpha}$ and $\alpha\delta > k(\omega_B, \gamma)$: Policy A *will* be chosen iff $\gamma\omega_B - 0.135 > \omega_B - 0.135$ or $\gamma > 1$, which is always true (by assumption).

These results are summarized in the following observation.

Observation:

1) When $\alpha < \hat{\alpha}$, Policy A will be chosen iff $\delta > \hat{\delta}_1 = \frac{1}{2\alpha} ln\left(\frac{\gamma}{(\gamma-1)+e^{-2\alpha}}\right)$.
2) When $\alpha > \hat{\alpha}$, Policy A will be chosen iff $\delta > \hat{\delta}_2 = \frac{1}{2\alpha} ln\left(\frac{\gamma\omega_B}{(\gamma-1)\omega_B+e^{-2}}\right)$.

Note that: $\lim_{\alpha\to 0} \hat{\delta}_1 = \lim_{\alpha\to 0}\left(\frac{1}{2\alpha} ln\left(\frac{\gamma}{(\gamma-1)+e^{-2\alpha}}\right)\right) = \frac{1}{\gamma}$; $\frac{d\hat{\delta}_1}{d\alpha} < 0$ and $\frac{d\hat{\delta}_2}{d\alpha} < 0$; $\hat{\delta}_1(\hat{\alpha}) = \hat{\delta}_2(\hat{\alpha})$; and $\lim_{\alpha\to\infty} \hat{\delta}_2 = 0$.

A2. Section 3.1.3: $v(\omega) = \beta \ln \omega$; Wealth-Adjusted WTP

Focusing on a policy with benefit $\Delta u_B = \alpha$, the WTP of the median-wealth individual is:

$$m(\varpi) = \begin{cases} \varpi - 0.135 & , \quad \frac{1}{2}\ln\varpi + 1 \le \alpha \\ \varpi(1-e^{-2\alpha}) & , \quad \frac{1}{2}\ln\varpi + 1 > \alpha \end{cases}$$

Substituting $\varpi = \frac{1}{2}(1+\gamma)\omega_B$, we obtain:

$$\frac{m(\varpi)}{m_B} = \begin{cases} \frac{1}{2}(1+\gamma) & , \quad \alpha \le \frac{1}{2}\ln\omega_B + 1 \\ \frac{\varpi(1-e^{-2\alpha})}{\omega_B - 0.135} & , \quad \alpha \in \left(\frac{1}{2}\ln\omega_B + 1, \frac{1}{2}\ln\varpi + 1\right) \\ \frac{\frac{1}{2}(1+\gamma)\omega_B - 0.135}{\omega_B - 0.135} & , \quad \alpha \ge \frac{1}{2}\ln\varpi + 1 \end{cases}$$

A3. Section 3.1.2: $v(\omega) = \sqrt{\omega}$; Individualized WTP

Substituting $v(\omega) = \sqrt{\omega}$ into the Equation (1), we have: $\sqrt{\omega_i} - \sqrt{\omega_i - m_i} = \Delta u_i$. Solving for m_i, we obtain:

$$m_i = min\left(\Delta u_i\left(2\sqrt{\omega_i} - \Delta u_i\right), \omega_i\right)$$

The *min* operator captures the wealth constraint: The individual would never be able to pay more than ω_i. (Note: In this example, we assume that the WTP is constrained only by the individual's wealth, ω_i, such that the individual can be left with zero, in contrast with the subsistence level assumption from the previous example.)

We have two ranges:

1) If $\sqrt{\omega_i} \leq \Delta u_i$ (or $\omega_i \leq \Delta u_i^{\,2}$), then $m_i = \omega_i$.
   If the utility from the policy change exceeds the utility obtained if the individual's wealth is spent entirely on the numeraire good, the individual would be willing to pay her entire wealth for the policy change.
   [When wealth is small relative to the magnitude of the policy change, i.e., when $\sqrt{\omega_i} \leq \Delta u_i$ (or $\omega_i \leq \Delta u_i^{\,2}$), the effect of wealth on WTP is: $\frac{dm_i}{d\omega_i} = 1 > 0$.]
2) If $\sqrt{\omega_i} > \Delta u_i$ (or $\omega_i > \Delta u_i^{\,2}$), then $m_i = \Delta u_i\left(2\sqrt{\omega_i} - \Delta u_i\right)$.
   If the utility from the policy change is smaller than the utility obtained if the individual's wealth is spent entirely on the numeraire good, the individual would be willing to pay $\Delta u_i\left(2\sqrt{\omega_i} - \Delta u_i\right) < \omega_i$ for the policy change.
   [When wealth is larger relative to the magnitude of the policy change, i.e., when $\sqrt{\omega_i} > \Delta u_i$ (or $\omega_i > \Delta u_i^{\,2}$), the effect of wealth on WTP is: $\frac{dm_i}{d\omega_i} = \frac{\Delta u_i}{\sqrt{\omega_i}} > 0$. ($\frac{dm_i}{d\omega_i} < 1$, but always positive; although the effect of wealth is increasing at a decreasing rate.)]

We can now consider the informational content of WTP. Starting with the first measure of informational content, $\frac{dm_i}{d\Delta u_i}$, we observe that:

1) When wealth is small relative to the magnitude of the policy change, i.e., when $\sqrt{\omega_i} \leq \Delta u_i$ (or $\omega_i \leq \Delta u_i^{\,2}$), WTP is a very poor proxy for utility. In fact, WTP carries no information about utility: $\frac{dm_i}{d\Delta u_i} = 0$.
2) When wealth is larger relative to the magnitude of the policy change, i.e., when $\sqrt{\omega_i} > \Delta u_i$ (or $\omega_i > \Delta u_i^{\,2}$), the informational content of WTP is: $\frac{dm_i}{d\Delta u_i} = 2\left(\sqrt{\omega_i} - \Delta u_i\right) > 0$. The informational content of WTP is increasing in wealth. When a person is richer, the utility that she obtains from the last units of the numeraire good is small. Therefore, for any increase in the value of the policy change, she would be willing to give up more units of the numeraire good.[41]

Next, consider the second measure of informational content. Consider Policy A that affects Individual A and Policy B that affects Individual B. Policy A increases Individual A's utility by Δu_A and Policy B increases Individual B's utility by Δu_B. If Individual A and Individual B have the same wealth, $\omega_A = \omega_B$, then A's WTP will exceed B's WTP, i.e., $m_A > m_B$ iff $\Delta u_A > \Delta u_B$. However, the greater the wealth disparity, the more likely it is that $m_A > m_B$ even though $\Delta u_A < \Delta u_B$.

We next attempt to quantify this distortion. We assume, without loss of generality, that A is richer than B and let $\omega_A = \gamma\omega_B$, with $\gamma > 1$. We measure the utilities from the policy change in relation

[41] Similar results if we use elasticity to measure the informational content of WTP:

$$\frac{dm_i^1/m_i^1}{d\Delta u_i/\Delta u_i} = \frac{dm_i^1}{d\Delta u_i}\cdot\frac{\Delta u_i}{m_i^1} = 2\left(\sqrt{\omega_i} - \Delta u_i\right)\cdot\frac{\Delta u_i}{\Delta u_i\left(2\sqrt{\omega_i} - \Delta u_i\right)} = \frac{2\sqrt{\omega_i} - 2\Delta u_i}{2\sqrt{\omega_i} - \Delta u_i}$$

The informational content of WTP is increasing in the individual's wealth.

to the utility that Individual B would have received had she spent her entire wealth on the numeraire good. Specifically, we have $\Delta u_B = \alpha v(\omega_B) = \alpha\sqrt{\omega_B}$, where $\alpha \geq 0$. And, for Individual A, we have $\Delta u_A = \delta \Delta u_B$, where $\delta \geq 0$. Or, substituting $\Delta u_B = \alpha\sqrt{\omega_B}$, we have $\Delta u_A = \delta\alpha\sqrt{\omega_B}$.

Now consider the parties' WTP. For Individual B, we have

$$m_B = \begin{cases} \omega_B & , \quad \sqrt{\omega_B} \leq \Delta u_B \\ \Delta u_B\left(2\sqrt{\omega_B} - \Delta u_B\right) & , \quad \sqrt{\omega_B} > \Delta u_B \end{cases}$$

And, after substituting $\Delta u_B = \alpha\sqrt{\omega_B}$:

$$m_B = \begin{cases} \omega_B & , \quad \alpha \geq 1 \\ \alpha(2-\alpha)\omega_B & , \quad \alpha < 1 \end{cases}$$

And for Individual A, we have

$$m_A = \begin{cases} \omega_A & , \quad \sqrt{\omega_A} \leq \Delta u_A \\ \Delta u_A\left(2\sqrt{\omega_A} - \Delta u_A\right) & , \quad \sqrt{\omega_A} > \Delta u_A \end{cases}$$

And, after substituting $\omega_A = \gamma\omega_B$ and $\Delta u_A = \delta\alpha\sqrt{\omega_B}$:

$$m_A = \begin{cases} \gamma\omega_B & , \quad \alpha \geq \dfrac{\sqrt{\gamma}}{\delta} \\ \delta\alpha\left(2\sqrt{\gamma} - \delta\alpha\right)\omega_B & , \quad \alpha < \dfrac{\sqrt{\gamma}}{\delta} \end{cases}$$

Policy A *will* be chosen iff $m_A > m_B$.

It is helpful to distinguish between two scenarios based on the magnitude of the policy change.

Scenario 1: Smaller Policy Changes, i.e., $\alpha < 1$
There are three ranges:

(i) $\delta < \sqrt{\gamma} < \frac{\sqrt{\gamma}}{\alpha}$ [or $\alpha < 1 < \frac{\sqrt{\gamma}}{\delta}$]: Policy A *will* be chosen iff $\delta\alpha(2\sqrt{\gamma} - \delta\alpha) > \alpha(2-\alpha)$

(ii) $\sqrt{\gamma} < \delta < \frac{\sqrt{\gamma}}{\alpha}$ [or $\alpha < \frac{\sqrt{\gamma}}{\delta} < 1$]: Policy A *will* be chosen iff $\delta\alpha(2\sqrt{\gamma} - \delta\alpha) > \alpha(2-\alpha)$

(iii) $\sqrt{\gamma} < \frac{\sqrt{\gamma}}{\alpha} < \delta$ [or $\frac{\sqrt{\gamma}}{\delta} < \alpha < 1$]: Policy A *will* be chosen iff $\gamma > \alpha(2-\alpha)$, which is always true (since $\alpha(2-\alpha) < 1$ for all α, and $\gamma > 1$ by assumption).

Scenario 2: Larger Policy Changes, i.e., $\alpha > 1$
There are three ranges:

(i) $\delta < \frac{\sqrt{\gamma}}{\alpha} < \sqrt{\gamma}$ [or $1 < \alpha < \frac{\sqrt{\gamma}}{\delta}$]: Policy A *will* be chosen iff $\delta\alpha(2\sqrt{\gamma} - \delta\alpha) > 1$

(ii) $\frac{\sqrt{\gamma}}{\alpha} < \delta < \sqrt{\gamma}$ [or $1 < \frac{\sqrt{\gamma}}{\delta} < \alpha$]: Policy A *will* be chosen iff $\gamma\omega_B > \omega_B$ or $\gamma > 1$, which is always true (by assumption).

(iii) $\frac{\sqrt{\gamma}}{\alpha} < \sqrt{\gamma} < \delta$ [or $\frac{\sqrt{\gamma}}{\delta} < 1 < \alpha$]: Policy A *will* be chosen iff $\gamma\omega_B > \omega_B$ or $\gamma > 1$, which is always true (by assumption).

In both scenarios, there exists a threshold value $\hat{\delta}$, such that Policy A will be chosen when $\delta > \hat{\delta}$ and Policy B will be chosen when $\delta < \hat{\delta}$. In Scenario 1, the threshold, which will be denoted $\hat{\delta}_1$, solves the equation $\hat{\delta}_1\alpha\left(2\sqrt{\gamma} - \hat{\delta}_1\alpha\right) = \alpha(2-\alpha)$, and we get: $\hat{\delta}_1 = \frac{\sqrt{\gamma}}{\alpha} - \frac{\sqrt{\gamma-\alpha(2-\alpha)}}{\alpha}$. (The equation that defines the threshold $\hat{\delta}_1$ appears in Ranges (i) and (ii); in Range (iii), δ is always above $\hat{\delta}_1$ and Policy A will always be chosen.) In Scenario 2, the threshold, which will be denoted $\hat{\delta}_2$, solves the equation $\hat{\delta}_2\alpha\left(2\sqrt{\gamma} - \hat{\delta}_2\alpha\right) = 1$, and we get: $\hat{\delta}_2 = \frac{\sqrt{\gamma}}{\alpha} - \frac{\sqrt{\gamma-1}}{\alpha}$. (The equation that defines the threshold $\hat{\delta}_2$ appears in Range (i); in Ranges (ii) and (iii), δ is always above $\hat{\delta}_2$ and Policy A will always be chosen.)

These results are summarized in the following observation.

Observation:

1) In Scenario 1, Policy A will be chosen iff $\delta > \hat{\delta}_1 = \frac{\sqrt{\gamma}}{\alpha} - \frac{\sqrt{\gamma-\alpha(2-\alpha)}}{\alpha}$.
2) In Scenario 2, Policy A will be chosen iff $\delta > \hat{\delta}_2 = \frac{\sqrt{\gamma}}{\alpha} - \frac{\sqrt{\gamma-1}}{\alpha}$.

Note that: $\lim_{\alpha\to 0} \hat{\delta}_1 = \lim_{\alpha\to 0}\left(\frac{\sqrt{\gamma}}{\alpha} - \frac{\sqrt{\gamma-\alpha(2-\alpha)}}{\alpha}\right) = \frac{\sqrt{\gamma}}{\gamma}$; $\frac{d\hat{\delta}_1}{d\alpha} < 0$ and $\frac{d\hat{\delta}_2}{d\alpha} < 0$ [CONFIRM]; $\hat{\delta}_1(\alpha = 1) = \hat{\delta}_2(\alpha = 1) = \sqrt{\gamma} - \sqrt{\gamma - 1}$; and $\lim_{\alpha\to\infty} \hat{\delta}_2 = 0$.

We derived a closed-form expression for the threshold $\hat{\delta}$:

$$\hat{\delta} = \begin{cases} \hat{\delta}_1 = \dfrac{\sqrt{\gamma}}{\alpha} - \dfrac{\sqrt{\gamma - \alpha(2-\alpha)}}{\alpha} & , \quad \alpha < 1 \\ \hat{\delta}_2 = \dfrac{\sqrt{\gamma}}{\alpha} - \dfrac{\sqrt{\gamma - 1}}{\alpha} & , \quad \alpha \geq 1 \end{cases}$$

In Figure A1 (below), we graph the threshold $\hat{\delta}$ as a function of α (the solid black line).

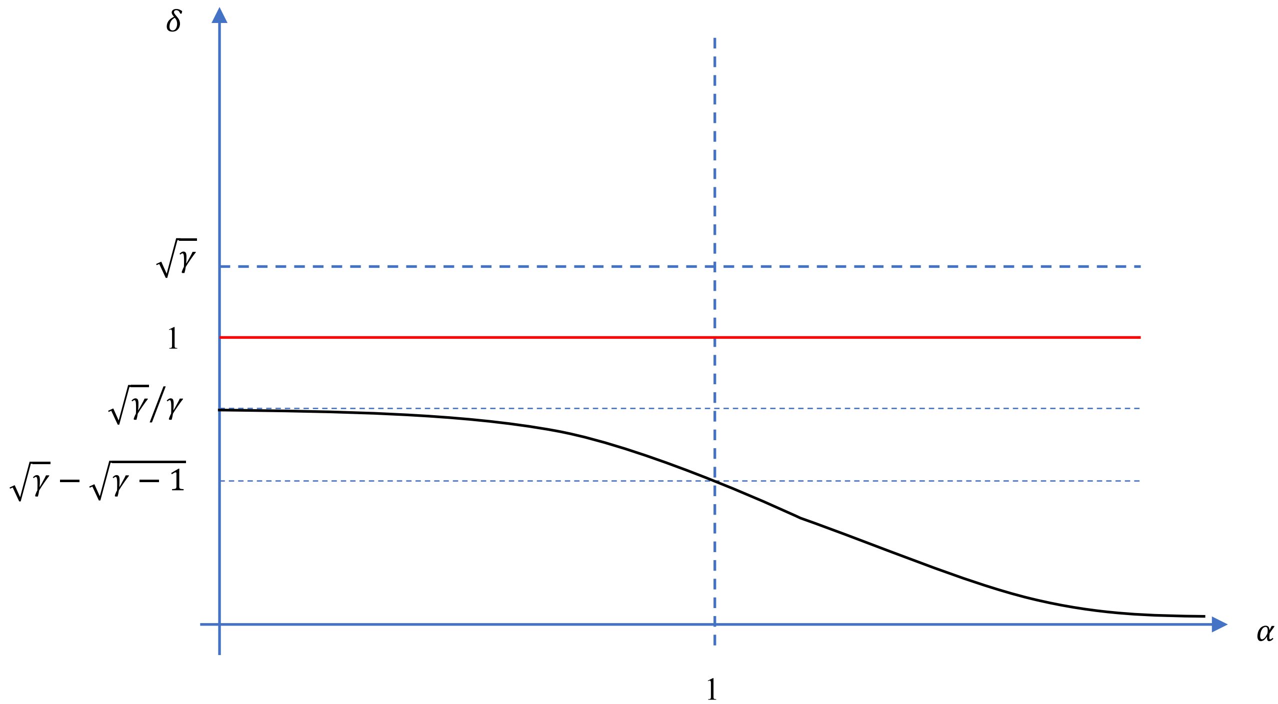

Figure A1: The threshold $\hat{\delta}$ as a function of α; $v(\omega) = \sqrt{\omega}$

Policy A *should* be chosen in the area above the red line, i.e., when $\delta > 1$. With WTP-based policymaking, Policy A *will* be chosen in the area above the black line, i.e., when $\delta > \hat{\delta}_1$ (if $\alpha < 1$) or $\delta > \hat{\delta}_2$ (if $\alpha > 1$). The distortion zone is the area between the black and red lines. The size of this zone increases with the magnitude of the wealth disparity (γ). The size of the distortion is also increasing in the impact of the policy (α). With a bigger policy change, the effect of wealth is larger.[42] And the distortion caused by wealth disparity is correspondingly larger.

Some simple comparative statics can help to ascertain the size of the distortion that is caused by wealth disparity. Starting with the magnitude of the wealth disparity, γ: What is a plausible value for this parameter? Since the simple example includes only two representative individuals, A and B, with different levels of wealth, let's divide the US population into two groups – the top 50% and the bottom 50% in terms of wealth (net worth). The median household in the "top" group, the household in the 75th percentile, has wealth of $\omega_A = \sim\$400{,}000$. The median household in the "bottom" group, the household in the 25th percentile, has wealth of $\omega_B = \sim\$10{,}000$. [Based on Federal Reserve estimates from 2016. GET more recent figures.] This implies: $\gamma = 40$. We thus have: $\hat{\delta}_1 = \frac{\sqrt{40}}{\alpha} - \frac{\sqrt{40-\alpha(2-\alpha)}}{\alpha}$ and $\hat{\delta}_2 = \frac{\sqrt{40}}{\alpha} - \frac{\sqrt{39}}{\alpha}$. Let's consider policy changes with different impact levels, i.e., with different α values. The results are summarized in the following table.

[42] Compare: WTP for a house is more sensitive to wealth than WTP for a pencil. At the extreme, we fall in the $v(\omega_i) < \Delta u_i$ range, where the individual would use all of her wealth to pay for the policy; and then WTP is determined solely by the individual's wealth.

α	$\hat{\delta}$	$1-\hat{\delta}$	$1/\hat{\delta}$
0.01	0.16	0.84	6.25
0.25	0.14	0.86	7.14
0.5	0.12	0.88	8.33
1	0.08	0.92	12.5
1.5	0.05	0.95	20
2	0.04	0.96	25

Table A1: $\gamma = 40$

Focusing on the fourth column in Table 1, we see that for a smaller-impact policy, e.g., $\alpha = 0.01$, a WTP-based assessment will prescribe Policy A even if it produces utility that is 6.25 times smaller than Policy B; and for a larger-impact policy, e.g., $\alpha = 2$, a WTP-based assessment will prescribe Policy A even if it produces utility that is 25 times smaller than Policy B.

What if we had a much smaller wealth disparity, represented by $\gamma = 4$ (rather than $\gamma = 40$)? We would have: $\hat{\delta}_1 = \frac{\sqrt{4}}{\alpha} - \frac{\sqrt{4-\alpha(2-\alpha)}}{\alpha}$ and $\hat{\delta}_2 = \frac{\sqrt{4}}{\alpha} - \frac{\sqrt{3}}{\alpha}$. Table 2 reports comparative statics for policy changes with different impact levels, i.e., with different α values.

α	$\hat{\delta}$	$1-\hat{\delta}$	$1/\hat{\delta}$
0.01	0.5	0.5	2
0.25	0.45	0.55	2.22
0.5	0.39	0.61	2.56
1	0.27	0.73	3.7
1.5	0.18	0.82	5.56
2	0.13	0.87	7.69

Table A2: $\gamma = 4$

Focusing on the fourth column in Table 2, we see that for a smaller-impact policy, e.g., $\alpha = 0.01$, a WTP-based assessment will prescribe Policy A even if it produces utility that is two times smaller than Policy B; and for a larger-impact policy, e.g., $\alpha = 2$, a WTP-based assessment will prescribe Policy A even if it produces utility that is 7.69 times smaller than Policy B.

A4. Section 4.2.2: $v(\omega) = \beta \ln \omega$; Uniform, Average WTP

For Individual B, we have

$$m_B = \begin{cases} \omega_B - 0.135 & , \quad \frac{1}{2}\ln \omega_B + 1 \leq \Delta u \\ \omega_B(1 - e^{-2\Delta u}) & , \quad \frac{1}{2}\ln \omega_B + 1 > \Delta u \end{cases}$$

And, after substituting $\Delta u = \alpha$:

$$m_B = \begin{cases} \omega_B - 0.135 & , \quad \alpha \geq \hat{\alpha}_B \\ \omega_B(1 - e^{-2\alpha}) & , \quad \alpha < \hat{\alpha}_B \end{cases}$$

where $\hat{\alpha}_B \equiv 1 + \frac{1}{2}\ln \omega_B$. (Note that $\omega_B > 0.135$ implies $\hat{\alpha}_B > 0$.)

And for Individual A, we have

$$m_A = \begin{cases} \omega_A - 0.135 & , \quad \dfrac{1}{2}\ln \omega_A + 1 \leq \Delta u \\ \omega_A(1 - e^{-2\Delta u}) & , \quad \dfrac{1}{2}\ln \omega_A + 1 > \Delta u \end{cases}$$

And, after substituting $\Delta u = \alpha$ and $\omega_A = \gamma\omega_B$:

$$m_A = \begin{cases} \gamma\omega_B - 0.135 & , \quad \alpha \geq \hat{\alpha}_A \\ \gamma\omega_B(1 - e^{-2\alpha}) & , \quad \alpha < \hat{\alpha}_A \end{cases}$$

where $\hat{\alpha}_A \equiv 1 + \frac{1}{2}\ln \gamma\omega_B$. (Note that $\omega_B > 0.135$ and $\gamma > 1$ imply $\hat{\alpha}_A > 0$.)

$$\bar{m} = \frac{1}{2}(m_A + m_B) = \begin{cases} \dfrac{1}{2}(1+\gamma)\omega_B - 0.135 & , \quad \alpha \geq \hat{\alpha}_A \\ \dfrac{1}{2}[\omega_B - 0.135 + \gamma\omega_B(1 - e^{-2\alpha})] & , \quad \hat{\alpha}_B \leq \alpha < \hat{\alpha}_A \\ \dfrac{1}{2}(1+\gamma)\omega_B(1 - e^{-2\alpha}) & , \quad \alpha < \hat{\alpha}_B \end{cases}$$