ECONSTOR
Make Your Publications Visible.

A Service of

ZBW Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Capra, C. Monica; Kniesner, Thomas J.

Article — Published Version

Daniel Kahneman's underappreciated last published paper: Empirical implications for benefit-cost analysis and a chat session discussion with bots

Journal of Risk and Uncertainty

Provided in Cooperation with:
Springer Nature

Suggested Citation: Capra, C. Monica; Kniesner, Thomas J. (2025) : Daniel Kahneman's underappreciated last published paper: Empirical implications for benefit-cost analysis and a chat session discussion with bots, Journal of Risk and Uncertainty, ISSN 1573-0476, Springer US, New York, NY, Vol. 71, Iss. 1, pp. 29-51, https://doi.org/10.1007/s11166-025-09459-5

This Version is available at:
https://hdl.handle.net/10419/330741

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CC BY https://creativecommons.org/licenses/by/4.0/

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

Journal of Risk and Uncertainty (2025) 71:29–51
https://doi.org/10.1007/s11166-025-09459-5

Daniel Kahneman's underappreciated last published paper: Empirical implications for benefit-cost analysis and a chat session discussion with bots

C. Monica Capra[1,2] · **Thomas J. Kniesner**[3,4,5]

Accepted: 29 April 2025 / Published online: 7 June 2025
© The Author(s) 2025

Abstract

Nobel Prize winner Daniel Kahneman's last published paper is an adversarial collaboration in which he and Matthew Killingsworth reconcile conflicting empirical results from their previous research on income and reported happiness, with Barbara Mellers as a facilitator. The empirical results use quantile regression to allow for measured income heterogeneity effects that include notch points in the estimated marginal utilities of income. Our analysis examines Kahneman's last paper's conceptual innovations and challenges to assumptions about diminishing marginal utility of income. We review his contributions to emotional well-being measurement and employ a novel AI-simulated dialogue between the late Amos Tversky and Sir Angus Deaton to explore interdisciplinary perspectives on the findings. Our paper demonstrates how Kahneman's final research undermines recent arguments for incorporating income redistribution simply into benefit-cost analysis, suggesting that such objectives remain better addressed through fiscal policy rather than regulatory interventions. His final published work exemplifies Kahneman's commitment to empirical precision and theoretical flexibility, even when contradicting his earlier conclusions.

Keywords Income satiation · Well-being · Adversarial collaboration · Quantile regression · Marginal utility · Social welfare weights · Simulated dialogue with AI

✉ Thomas J. Kniesner
tom.kniesner@cgu.edu

[1] Claremont Graduate University, Claremont, CA, USA

[2] Center for the Philosophy of Freedom, University of Arizona, Tucson, AZ, USA

[3] Economics Emeritus, Syracuse University, Syracuse, NY, USA

[4] Institute of Labor Economics, IZA, Bonn, Germany

[5] Department of Economic Sciences, Claremont Graduate University, Claremont, CA, USA

JEL codes D12 · D61 · H23 · I31

1 Introduction

We are honored to examine the final research contribution of the late Nobel Prize Winner in Economic Sciences, Daniel Kahneman, as presented in what we contend is an underappreciated last published paper (Killingsworth et al., 2023). Kahneman's last published paper is a conceptually sophisticated adversarial collaboration reconciling conflicting joint empirical patterns of income and experienced well-being that is important for public policy having distributional consequences.

Killingsworth's (2021) investigation suggested that happiness increases linearly with the logarithm of income for everyone. In contrast, the earlier research of Kahneman and Deaton (2010) indicated a flattening pattern where happiness plateaus beyond a certain income level. The core empirical tension between the two sets of findings involves two fundamental questions concerning marginal utility. First, "Does reported well-being continue to rise at incomes above some threshold?" The answer is "Yes." Second, "Is the estimated marginal utility of income at relatively high incomes lower than at relatively low incomes?" Contrary to conventional economic assumptions, the data suggest that the answer is "No." Through meticulous reanalysis of Killingsworth's (2021) experience sampling data, the adversarial collaboration involving Kahneman in Killingsworth et al. (2023) discovered (1) that the flattening pattern of utility with income was primarily observed in the least happy 20% of the population, and (2) that happiness continued to increase with income for the happier majority and exhibited accelerating returns for individuals in the highest happiness quantiles.

The empirical pattern fleshed out by Kahneman and co-authors in Killingsworth et al. (2023) yields a significant refinement in our understanding of the income-happiness relationship. Kahneman's evolution from his earlier work with Deaton in Kahneman and Deaton (2010) on utility and income demonstrated a commitment to empirical precision even when it contradicted his previous findings, allowing data-driven insights to supersede theoretical expectations. By employing quantile regression techniques, the Killingsworth et al. (2023) collaboration revealed systematic heterogeneity in well-being responses to income that traditional analytical frameworks had obscured. The findings challenge homogeneity assumptions that have long underpinned economic utility and welfare assessment models. Kahneman's research trajectory exemplifies how refinements in the method can fundamentally alter substantive conclusions about the relationship between income and happiness.

Here, we systematically examine Kahneman's contributions to happiness research through multiple complementary dimensions. We first analyze his conceptual innovations in measuring experienced well-being, highlighting how technical refinements in assessment procedures captured previously undetectable patterns in subjective experience. We then explore how his research illuminated complex determinants of happiness, with particular attention to an evolving model of income effects across different segments of the well-being distribution. Finally, we consider the substantive policy implications emerging from Kahneman's research trajectory, specifically how the

revised empirical patterns identified in his final publication necessitate reconsideration of economic interventions designed to enhance population well-being. By synthesizing Kahneman's theoretical contributions and empirical insights, we provide a comprehensive assessment of his enduring legacy in studying happiness and subjective experience. His final collaboration exemplifies a scientific approach characterized by empirical rigor, theoretical flexibility, and willingness to revise established conclusions in response to emerging evidence—qualities that defined his transformative impact on behavioral economics and psychological science.

In exploring the findings on income and happiness, we incorporate a novel element: an AI-simulated chatbot discussion featuring two of Kahneman's most influential collaborators: his longtime collaborator, the late Amos Tversky, and Sir Angus Deaton, who co-authored their 2010 study. Although neither Tversky nor Deaton participated in the adversarial collaboration we focus on here, both played pivotal roles in shaping Kahneman's understanding of income and well-being.

Recent studies show that large language models (LLMs) can be useful for simulating decision-makers (Horton, 2023; Anthis et al., 2025; Gonzalez-Bonorino, Capra, and Pantoja, 2025). In Gonzalez-Bonorino et al. (2025), we profiled synthetic cultural agents using retrieval-augmented generation (RAG), but prompt engineering and other tools have also been successfully used (see Anthis et al., 2025). Following their previous research, we use LLMs to profile Tversky and Deaton and prompt the two synthetic agents to engage in dialogue regarding the adversarial collaboration of Killingsworth, Kahneman, and Meller. The AI-simulated dialogue presented here follows Kahneman's innovative approach to social science research. The dialogue between profiled synthetic agents called Amos and Angus also offers additional insights into the logic and conclusions of Killingsworth et al. (2023).

Finally, we connect the underappreciated results of Killingsworth et al. (2023) to the ongoing debate about whether distributional considerations should be incorporated into benefit-cost analysis (Viscusi, 2024). Advocates for expanding the boundaries of the benefit-cost approach argue that measuring social utility enhancing redistributive effects of transferring income from wealthier to poorer individuals can lead to more policies to enhance social welfare. However, the findings of Killingsworth et al. (2023) challenge this. Their results suggest that the diminishing marginal utility of income needed to justify redistribution within benefit-cost frameworks, is not empirically supported. Instead, their paper reinforces the argument that income redistribution is better handled through fiscal policy, such as taxes and transfers, rather than through regulatory interventions.

The remainder of our paper is organized as follows. Section 2 reviews Kahneman's contributions to the measurement of subjective well-being and the findings of Kahneman's last published paper. Section 2 also provides insights from a simulated AI-based chatbot discussion featuring Amos Tversky and Sir Angus Deaton. Section 3 explores how the findings of Killingsworth et al. (2023) relate to the ongoing debate about incorporating distributional considerations into benefit-cost analysis. Finally, Sect. 4 concludes by reflecting on the broader implications of the findings in Kahneman's last published paper for future research and policy design.

2 The focus of Kahneman's last published paper

Kahneman's last published paper, Killingsworth et al. (2023), is relevant to economists because it highlights how almost contradictory conclusions can be drawn about the same question when different methods and similar assumptions are not carefully assessed. Killingsworth's 2021 study and Kahneman and Deaton's research differed significantly in their methods, primarily in how they measured emotional well-being and collected data. Yet, the research in both papers made implicit assumptions about the homogeneity of the relationship between income and well-being across the population.

2.1 Measuring subjective Well-Being: Kahneman's contributions and implications

Traditional measurements of well-being have primarily focused on evaluative assessments, which require respondents to make global judgments about their life circumstances. Assessments often use scales such as Cantril's ladder, where respondents rate their life satisfaction on a scale from 0 to 10 or other similar evaluative questions that capture a person's general feelings about their life circumstances (Stevenson & Wolfers, 2013; Kahneman & Krueger, 2006). For example, the Gallup World Poll measures happiness by asking respondents to rate their life satisfaction on a scale from 0 (worst possible life) to 10 (best possible life). Other widely used surveys, such as the General Social Survey (GSS), measure happiness by asking people the following question: "Taken all together, how would you say things are these days—would you say that you are very happy, pretty happy, or not too happy?"

Evaluative assessment approaches have been foundational in economic research, aligning closely with neoclassical utility theory, which posits that individuals make choices based on their preferences and the perceived utility derived from them (Bryson & MacKerron, 2016). However, as Kahneman et al. (1997) pointed out, the traditional methods rely on cognitive evaluations that may not accurately reflect individuals' emotional experiences. Kahneman saw the traditional approach of measuring happiness as a consequence of economists' emphasis on "decision utility," which refers to the value individuals assign to outcomes when making decisions and reflects the anticipation or remembered utility. Evaluative assessments are based on recollections of how enjoyable or painful an experience was or will be.

Kahneman observed that anticipatory and evaluative responses rely on cognitive evaluations that are subject to biases. For instance, the "focusing illusion" bias suggests that people tend to overemphasize certain aspects of their lives when evaluating their happiness (Kahneman, 2011). In the context of evaluating life satisfaction and its relationship to income, participants might disproportionately weigh their financial situation when reporting their experienced well-being, leading to inflated perceptions of happiness associated with higher income levels. The "availability heuristic" (Tversky & Kahneman, 1997) could lead individuals to recall more positive experiences tied to their incomes, particularly if they have recently made a significant purchase or received a raise, thus distorting their overall assessment of happiness. Such cognitive biases manifest in several measurement artifacts: respondents overweight peak and end experiences during retrospective evaluation (Fredrickson & Kahneman, 1997),

exhibit heightened sensitivity to contextual comparisons, and demonstrate duration neglect when integrating temporal experiences (Kahneman, 1997). Collectively, measurement limitations undermine the validity of traditional approaches to capturing experiential well-being.

In response to measurement challenges, Kahneman proposed an alternative measurement framework emphasizing "experienced utility," focusing on the actual feelings and emotions individuals experience in their daily lives rather than their retrospective evaluations of life satisfaction (Kahneman & Sugden, 2005). The Experience Sampling Method (ESM), also known as the Ecological Momentary Assessment (EMA), was developed to collect information on people's reported feelings in real time in their natural environments (Kahneman et al., 2004). In ESM studies, participants carry an electronic diary (a phone app) that beeps at random times during the day and asks respondents to describe what they were doing just before the prompt and indicate the intensity of various feelings (happy, frustrated/annoyed).

Although ESM is considered the gold standard for measuring well-being in the Edgeworth tradition because it minimizes the effects of judgment and recall bias, it has proven prohibitively expensive and burdensome to administer to large representative samples (Kahneman et al., 2004; Benjamin et al., 2012). Additionally, ESM has very low response rates for certain activities and rarely captures infrequent activities (Kahneman et al., 2004). Recognizing the practical difficulties, Kahneman subsequently developed the Day Reconstruction Method (DRM) as a practical alternative that could approximate measurement precision like the ESM while being easier to implement.

In the DRM, respondents fill out a diary corresponding to episodes of their previous day, describing when the episode began and ended, what they were doing, where they were, whom they were with, and how they felt during each episode, along with selected affect dimensions (Kahneman et al., 2004). Research has shown that the DRM successfully reproduces the patterns of emotional experience captured by ESM, including complex patterns like the diurnal variation in tiredness and positive/negative affect while being much more practical to implement on a large scale (Kahneman et al., 2004; Benjamin et al., 2012).

DRM reduces biases that can distort self-reported evaluations and, perhaps more importantly, the DRM framework distinguishes between evaluative well-being (global life judgments) and experienced well-being (moment-to-moment affective states). Although the DRM represents a significant advancement in measuring experienced utility, it is not without limitations. One key concern is that the Day Reconstruction Method relies on participants' ability to accurately recall and reconstruct their emotional states from the previous day. The reliance on memory introduces the potential for recall biases, such as the "peak-end rule," where individuals disproportionately weigh the most intense and final moments of an experience (Fredrickson & Kahneman, 1997). Additionally, the passage of time between the experience and its reconstruction may lead to distortions influenced by subsequent events or contextual factors. Despite its limitations, the DRM remains a valuable tool for approximating real-time emotional experiences, particularly when the logistical and financial constraints of administering the Experience Sampling Method to large, representative samples are prohibitive.

The evolution of the research method represents more than incremental refinement; it constitutes a paradigmatic shift in well-being assessment. By capturing the dynamic, contextual nature of emotional experience, the DRM reveals patterns that traditional evaluative measures systematically obscure and the advantages of Kahneman's proposed measurement approach are manifold. First, it provides a more holistic understanding of well-being by capturing the dynamic nature of emotional experiences, which can fluctuate throughout the day (Kahneman et al., 2004). Real-time assessment allows researchers to identify specific activities or contexts that contribute to higher or lower levels of happiness, offering insights that traditional measures may overlook. Second, by focusing on experienced utility, Kahneman's method addresses the limitations of cognitive biases inherent in retrospective evaluations, such as the focusing illusion, where individuals overestimate the impact of certain events on their overall happiness (Ubel et al., 2005).

The conceptualization divergence between Killingsworth (2021) and Kahneman and Deaton (2010) represents a critical explanatory factor in their apparently contradictory findings regarding income and happiness. Mathew Killingsworth used the experience sampling method, where participants were prompted on their smartphones to report their current happiness in real time. They answered the question, "How do you feel right now?" on a continuous scale, multiple times a day over several weeks. In contrast, Kahneman and Deaton (2010) used data from the Gallup-Healthways Well-Being Index, which included dichotomous (yes/no) questions about the emotional experience of the preceding day. The Kahneman/Deaton approach is less aligned with Kahneman's own work, as participants were asked to recall whether they experienced specific feelings such as happiness, enjoyment, and sadness, providing a less granular and more general recall of the previous day's emotions.[1]

The collaboration of Killingsworth, Kahneman, and Mellers re-examined data from Killingsworth (2021). With its continuous measurement scale, Killingsworth's original study showed a linear relationship between average experienced happiness and log(income) without a plateau.[2] The Killingsworth (2021) measurement approach is considered the "gold standard" for measuring emotional well-being, offering a more sensitive measure than the dichotomous (yes/no) questions used in the Kahneman and Deaton (2010) study. In fact, the dichotomous measures used by Kahneman and Deaton (2010) created a ceiling effect that distorted the observed relationship between income and well-being. For high-income respondents, who mostly reported positive affect, the measures could not differentiate between degrees of happiness. In contrast, the measures effectively captured variations in unhappiness by distinguish-

[1] The Gallup-Healthways Well-Being Index (GHWBI) and the Daily Reconstruction Method (DRM) are distinct approaches to measuring subjective well-being, differing in data collection and granularity. GHWBI employs a daily survey assessing the presence of certain feelings "during a lot of the day," relying on a more general recall of the previous day's emotions. DRM, however, uses a structured recall technique, where participants reconstruct their previous day by dividing it into episodes and reporting their feelings during each episode. By facilitating a more detailed reconstruction of the day's events, the DRM allows for more granular reporting of feelings tied to specific activities. In essence, while the DRM captures feelings with greater episodic detail, the GHWBI relies on a broader assessment of overall emotional experience.

[2] Killingsworth (2021) obtained "1,725,994 experience-sampling reports from 33,391 employed US adults."

ing those experiencing negative emotions from those not. Thus, the Kahneman and Deaton measures were better at detecting the absence of misery than the presence of happiness. Mistakenly interpreting their dependent variable as "happiness" instead of "unhappiness," Kahneman and Deaton overstated the extent to which happiness flattens above a certain income threshold.

2.2 Multidimensional determinants of subjective Well-Being

Kahneman's conceptual innovations have fundamentally transformed our understanding of the determinants of subjective well-being by revealing complex patterns that traditional evaluative measures systematically obscure. Traditional evaluative measures often emphasized economic factors, such as income and employment status, as primary determinants of happiness (Stevenson & Wolfers, 2013; Clark et al., 2008). However, Kahneman's focus on experienced utility has illuminated the importance of daily experiences, social interactions, and emotional states in shaping well-being. For instance, research utilizing the DRM has shown that individuals derive greater happiness from engaging in meaningful activities, such as spending time with loved ones, than from material wealth (Kahneman et al., 2004).

The shift in focus toward experienced utility and a moment-to-moment measurement of positive emotions has led to a more holistic view of what determines happiness. We now recognize that emotional experiences and social contexts are critical components of happiness. Well-being is influenced by factors such as social relationships and personal fulfillment (Easterlin, 2003; Krueger et al., 2009). For example, persons who report strong social connections tend to experience higher happiness and life satisfaction (Steptoe et al., 2015). Benjamin et al. (2012) found high relative marginal utilities for aspects related to family (well-being, happiness, and relationship quality), security (financial, physical, life, and the future in general), values (morality and meaning), and having options (freedom of choice and resources).

Health is another determinant that Kahneman's measures have highlighted. Studies have shown that individuals with better physical and mental health report higher levels of life satisfaction and emotional well-being (Steptoe et al., 2015). Kahneman's research indicates that health-related factors are integral to understanding subjective well-being, as they influence both the evaluative and experiential dimensions of happiness (Kahneman & Krueger, 2006; Kahneman & Riis, 2005). The interplay between health and well-being is particularly salient in aging populations, where maintaining physical health can lead to improved psychological outcomes (Steptoe et al., 2015).

Finally, Kahneman's research has shed light on the role of adaptation and social comparison in shaping people's well-being. People tend to adapt to changes in their circumstances over time, leading to a "hedonic treadmill" effect where increases in income or other factors have only a temporary impact on well-being (Brickman & Campbell, 1971; Kahneman, 1997). Additionally, people often compare their own circumstances to those of others, influencing their sense of well-being (Kahneman et al., 2006). Emotional well-being is also affected by social comparisons and the context in which individuals perceive their circumstances (Clark et al., 2008). Thus, the relationship between various determinants and subjective well-being is not always

straightforward and can be influenced by psychological processes such as adaptation and social comparison.

Kahneman's own research with Deaton (Kahneman & Deaton, 2010) indicates that although higher income improves life evaluation, it does not significantly enhance emotional well-being once basic needs are met. Their finding aligns with the notion of diminishing returns, where additional income beyond a certain threshold yields minimal increases in happiness. Integrating experienced utility into well-being research has prompted a reevaluation of the relative importance of various other determinants besides income. Daniel Kahneman's extensive work on the determinants of subjective well-being has highlighted the complexity of the factors influencing happiness and life satisfaction. His research has identified key determinants such as income, social relationships, health, personal values, context, and adaptation (Kahneman et al., 1997; Kahneman & Deaton, 2010).

Ironically, despite Kahneman's deep understanding of the complexities surrounding the determinants of subjective well-being, he and Deaton (2010) relied on the homogeneity assumption when analyzing the relationship between happiness and income. They plotted conditional means, which implicitly assumes that the mean values provide a sufficient description of how happiness relates to income across the entire population. Their approach overlooked the possibility that the distribution of happiness might change shape as income rises, failing to account for how the effects of income differed across the happiness distribution.

The collaborative reanalysis of Killingsworth's 2021 paper by Killingsworth, Kahneman, and Mellers in 2023 revealed that the relationship between experienced well-being and income was not consistent across the entire population and depended on an individual's position in the happiness distribution. The same increase in income has different effects on the happy and unhappy regions of the distribution. Although unhappiness diminishes with rising income for the relatively unhappy group up to a certain point, for the happier majority, happiness continues to rise with income and even accelerates in the happiest group.

Had Kahneman and Deaton (2010) incorporated the insights from Kahneman's own research on the determinants of subjective well-being, they might have anticipated that the relationship between income and happiness could vary across different subgroups of the population. A better understanding of the role of adaptation, social comparison, and individual circumstances would have highlighted the potential limitations of relying on the homogeneity assumption when studying the income-happiness relationship. Their oversight underscores the importance of carefully considering the complexity of subjective well-being and the need to question simplified assumptions, even when applied by experts in the field.

2.3 Public policy implications of multidimensional Well-Being assessment

The conceptual and empirical research innovations in Kahneman's subjective well-being research generate significant implications for public policy formulation and evaluation. As policymakers increasingly recognize well-being enhancement as a legitimate governmental objective beyond economic growth, integrating sophisti-

cated well-being measurement into policy processes becomes essential for evidence-based governance.

Behavioral economics research demonstrates that individuals frequently make choices that fail to maximize their experienced well-being, suggesting that policymakers should consider non-choice-based welfare indicators—particularly experiential well-being measures—when designing and evaluating interventions (Adler, 2016). By incorporating multidimensional well-being data into policy analysis, governments can develop more nuanced understanding of how regulatory and fiscal decisions influence citizens' actual lived experience rather than merely economic indicators.

The finding that income represents only one of multiple well-being determinants—and operates with variable impact across different population segments—suggests that policy portfolios should address social and emotional dimensions alongside economic factors. Initiatives fostering social connection, enhancing mental health service accessibility, and promoting community engagement may yield substantial well-being returns that purely economic interventions cannot achieve. Similarly, policies addressing income inequality and providing targeted support for disadvantaged populations may enhance subjective well-being across society through both absolute and relative improvement mechanisms (Stevenson & Wolfers, 2013).

Kahneman's work on measurement method further suggests that policy evaluation should incorporate multiple well-being assessment approaches rather than relying exclusively on global evaluative measures. By capturing both experiential and evaluative dimensions, policymakers can develop more comprehensive understanding of intervention impacts across different temporal scales and life domains. As demonstrated by Killingsworth et al. (2023), income affects well-being differently across the happiness distribution, with unique patterns for the least happy quintile that would be obscured in aggregate measures. The heterogeneous income-happiness relationship suggests specific policy implications: interventions should be tailored differently for those in the bottom quintile versus other population segments, potentially emphasizing non-economic factors for the least happy while pursuing different strategies for others. Such targeted approaches require measurement strategies sensitive to population differences and necessitate explicit consideration of potential tradeoffs between policies that improve average well-being versus those specifically designed to help the least happy. By maintaining this nuanced perspective on differential impacts, policymakers can develop more effective interventions recognizing human well-being's complex, multidimensional nature. Such a conceptual pluralism aligns with recent advances in public administration that emphasize multidimensional outcome measurement beyond traditional economic indicators.[3]

The 2023 collaboration's finding of heterogeneous income-happiness relationships across the well-being distribution has particularly significant policy implications. If happiness continues increasing with income for most population segments

[3] Krueger et al. (2009) provided an alternative to multidimensional measures. The U-index measures subjective well-being based on the proportion of time individuals spend in an unpleasant state; that is, moments when the most intense emotion experienced is negative, often provide more actionable insights for policy interventions than global life satisfaction, especially when focusing on improving the condition of the least happy populations.

but exhibits diminishing returns for the least happy quintile, targeted interventions addressing non-income determinants might be especially important for enhancing well-being among the most disadvantaged groups. This complex pattern suggests that policy design should consider not only average effects but also differential impacts across the well-being distribution.

2.4 Method innovation through adversarial collaboration

Kahneman's final research contribution exemplifies a conceptually sophisticated approach to scientific inquiry through adversarial collaboration—a structured framework for reconciling conflicting empirical findings through cooperative investigation by researchers with opposing views. Adversarial collaboration represents a distinctive innovative framework that transcends traditional research paradigms by directly addressing replication challenges that have emerged across psychological and economic sciences.

The adversarial collaboration of Kahneman, Killingsworth, and Mellers demonstrates several methodical virtues that enhance analytical rigor. First, it explicitly acknowledges theoretical tensions rather than marginalizing contradictory evidence. Second, it employs methodological pluralism by integrating multiple analytical techniques to examine the same underlying phenomena. Third, it introduces a neutral facilitator to mitigate confirmation biases that might otherwise distort interpretation. This structured approach to scientific disagreement provides a template for resolving empirical discrepancies that extends beyond subjective well-being research to the broader scientific enterprise.

Central to the collaboration's empirical contribution is its employment of distribution-sensitive analytical techniques—specifically quantile regression—that reveal systematic heterogeneity obscured by conventional mean-based approaches. By examining how income-happiness relationships vary across different segments of the well-being distribution, the research moves beyond simplistic questions about whether satiation occurs toward more nuanced examination of for whom and under what conditions income influences experienced well-being. The analytical refinement represents a significant advancement in how researchers conceptualize and investigate complex psychoeconomic relationships.

The revelation that the flattening pattern appears primarily among the least happy population quintile challenges conventional interpretations while simultaneously validating aspects of both original studies. The finding demonstrates that apparently contradictory results may reflect conceptual artifacts or unexamined distributional variations rather than fundamental incompatibility. By systematically investigating multiple alternative possibilities through collaborative analysis, the researchers exemplify scientific virtues of intellectual humility and empirical precision that characterize the most productive forms of academic inquiry.

2.5 Simulated academic discourse: Conceptual considerations and analytical insights

2.5.1 Theoretical rationale for simulated scholarly dialogue

The adversarial collaboration between Daniel Kahneman and Matthew Killingsworth, with Barbara Mellers as the facilitator/tiebreaker, naturally raises questions about how the specific method would be evaluated by those who shaped Kahneman's early contributions to behavioral economics. The absence of input from Kahneman's most significant intellectual collaborators, particularly Amos Tversky, who co-developed prospect theory and the focusing illusion concept, creates an interpretative gap in contextualizing the empirical results of interest here within the broader theoretical ecosystem of behavioral economics.

Tversky's pioneering work on decision-making heuristics provides especially relevant conceptual frameworks for studying the income-happiness relationship. For example, the focusing illusion—which systematically biases utility predictions through disproportionate attention to specific attributes—has direct implications for interpreting how individuals respond to income changes across different well-being segments. Similarly, Tversky's work on reference-dependent preferences suggests potential explanatory mechanisms for the observed heterogeneity in happiness responses to income across the distribution.

Although Tversky's absence from discourse since 1996 precludes direct engagement with the findings of interest here, advanced computational methods permit the simulation of theoretical perspectives based on established scholarly traditions. Simulations provide a counterfactual opportunity for economists: how might Tversky have critiqued the adversarial collaboration itself, particularly in terms of its approach to testing hypotheses about subjective well-being and utility? This is more than a historical curiosity for those working at the intersection of behavioral economics and welfare analysis. It offers a framework for critically examining whether the collaboration's methods adequately account for the systematic biases in utility estimation that Kahneman and Tversky originally identified, such as the risk of overemphasizing certain variables such as income.

Given that we cannot validate a simulated Tversky dialogue with the actual scholar, we expanded our approach to include Sir Angus Deaton, who is Kahneman's collaborator on the influential 2010 *PNAS* paper and a living economist whose expertise in welfare and well-being complements the discussion. Unlike Tversky, Deaton could potentially assess the accuracy of such simulations, providing a critical check on their validity. Our method of simulating a dialogue between Tversky and Deaton bridges past and present, ensuring that the simulated conversation remains both intellectually rigorous and directly relevant to evaluating collaborative research in economics.

2.5.2 Method and implementation

To further explore the relationship between income and subjective well-being, as discussed in the 2023 paper by the adversarial collaboration, we included a simulated dialogue between Amos Tversky and Angus Deaton. The dialogue was generated

using a custom-built simulation framework using Python. The code is found in our Online Appendix and can be accessed through the associated Link.[4]

The simulation framework employs LangChain to orchestrate the conversation flow and the Groq large language model to generate text responses. The personas of Amos Tversky and Angus Deaton were carefully constructed to reflect their respective expertise—Tversky as a cognitive psychologist specializing in heuristics and decision-making, and Deaton as an economist focused on poverty, welfare, and subjective well-being. The framework ensures that the dialogue remains grounded in the context of the 2023 paper, with prompts designed to elicit critical engagement with its findings, method, and implications.

Key components of the simulation include:

1. Character Profiles: Each persona features a detailed background and expertise, ensuring their responses align with their intellectual perspectives. Domain-specific knowledge graphs incorporate Tversky's prospect theory and Deaton's income inequality metrics.
2. Conversation Context: The simulation is guided by a specific topic, context, and key points, which include the findings of the 2023 paper and its reconciliation of earlier research on income and happiness.
3. Dialogue Generation: The conversation alternates between the two personas, with each turn informed by the dialogue history and the predefined context. Responses are structured to reflect the personas' expertise and to explore alternative explanations, methodological critiques, and broader implications. We adjusted the temperature and top-p to produce divergent argumentation patterns.
4. Error Handling and Formatting: The framework includes robust error handling to ensure smooth execution and readable output while enforcing strict formatting rules to maintain the integrity of the simulated dialogue.

The architecture enables the AI to generate responses that closely mimic each scholar's unique perspective and publication history, with Tversky favoring counterfactual reasoning and Deaton prioritizing evidentiary standards. The simulation framework demonstrates how AI can rigorously model contrasting scholarly viewpoints while maintaining the intellectual depth of human debate.

The full implementation of the simulation framework, including the code, can be accessed through the Online Appendix and this link[5] (Link) for transparency and reproducibility.

2.5.3 Insights on the dialogue output: Tversky vs. Deaton perspectives

The simulated dialogue revealed fundamental disciplinary tensions in approaching the income-happiness relationship. Tversky bot (Tversky bot) consistently framed the findings through the lens of prospect theory, particularly its emphasis on loss aversion. He proposed that the concave utility function for gains and the convex util-

[4] https://colab.research.google.com/drive/1Lt8tNzVWZzJckC3AMT3RBHE8Y6LCEmmt?usp=sharing.

[5] https://colab.research.google.com/drive/1Lt8tNzVWZzJckC3AMT3RBHE8Y6LCEmmt?usp=sharing.

ity function for losses could explain the income insensitivity observed in the unhappiest cohort. The psychological framing highlights the potential role of cognitive biases and heuristics in shaping subjective well-being. The Appendix at the end of our paper has a deeper discussion of the dialogue for the interested reader.

In contrast, Deaton bot (Deaton bot) approached the issue from an econometric perspective, emphasizing structural equation modeling. He conceptualized happiness as a function of the logarithm of income, alongside control variables, such as health and education, which collectively accounted for 38% of the variance in subjective well-being, as demonstrated in Kahneman and Deaton (2010). The method divergence underscores the complementary nature of cognitive psychology and applied microeconomics in understanding subjective well-being. Although Tversky's approach focuses on individual-level psychological mechanisms, Deaton's perspective emphasizes the importance of robust empirical modeling and measurement.

2.5.4 Bot interpretations on the relationship between income and happiness

The dialogue also provided novel interpretations of the 2023 findings on the relationship between income and happiness. Tversky bot proposed an adaptive expectations theory, suggesting that individuals with lower incomes develop psychological coping mechanisms that flatten perceived utility gains beyond $75,000 annually. This aligns with the idea that individuals adjust their expectations and emotional responses to their economic circumstances, potentially mitigating the impact of additional income on happiness.

Deaton bot, however, identified potential conceptual artifacts that could influence the observed patterns. He pointed to differential item functioning across income quartiles in the 15-item well-being scale, which could result in measurement inconsistencies. Additionally, he raised concerns about selection bias in experience sampling, noting that low-SES (socioeconomic status) participants often exhibit lower response rates in longitudinal studies. Deaton bot's critiques highlight the importance of addressing conceptualization challenges to ensure the validity of findings. Together, the contrasting perspectives illustrate the complexity of the happiness-income paradox and the value of interdisciplinary approaches in behavioral economics research.

2.5.5 Contributions and features of the simulation

Our study builds on Daniel Kahneman's interdisciplinary legacy by integrating psychological insights with economic analysis through computational dialogue. The framework demonstrates how AI systems can enhance behavioral economics research by simulating high-level academic discourse and bridging traditionally separate fields. The AI system facilitated a simulated academic debate, achieving a rebuttal density of 72% within two exchange cycles—indicating that 72% of the responses were counterarguments to prior points. The high level of engagement reflects the system's ability to replicate the dynamic nature of academic discourse.

The simulated Tversky persona employed constructs from prospect theory, such as loss aversion, to critically evaluate Deaton's econometric models of income thresholds. The process mirrors the intellectual tension between psychological mechanisms

and economic measurement that characterizes much of Kahneman's work. The model also achieved an 89% topical consistency, effectively aligning psychological constructs with economic metrics. For example, the focusing illusion—a psychological concept—was translated into the Gini coefficient, while loss aversion was connected to the income elasticity of happiness. The interdisciplinary translation reflects Kahneman's efforts to quantify subjective phenomena, uncovering new links between psychology and economics.

2.5.6 Output of simulated conversation

[Amos]: Let's discuss Killingsworth et al., 2023 *PNAS* paper entitled "Income and Emotional Well-Being: A Conflict Resolved." I'm particularly interested in how we can approach this from different perspectives. I'm intrigued by the findings that the flattening pattern observed by Kahneman and Deaton is primarily present in the least happy individuals. This raises questions about the ceiling effect in their measurement of happiness. However, I'm not convinced that focusing on unhappiness instead of happiness would have led to different conclusions in the original study. Can we explore alternative explanations for this finding, such as the role of expectations and adaptability in shaping our emotional well-being? For instance, do people with lower incomes have different expectations about their future prospects, which might influence their emotional well-being? Or is it possible that happier individuals are more likely to adapt to changes in their income, leading to continued increases in happiness? I'd love to hear Angus' thoughts on these possibilities.

[Angus]: Amos, I'm intrigued by the findings, but I remain skeptical about the conclusion that the flattening pattern is primarily present in the least happy individuals. I think we need to consider alternative explanations, such as the role of expectations and adaptability. For instance, do people with lower incomes have different expectations about their future prospects, which might influence their emotional well-being? Or is it possible that happier individuals are more likely to adapt to changes in their income, leading to continued increases in happiness? I'd like to see more evidence on how expectations and adaptability affect the relationship between income and emotional well-being. Additionally, I think we should also consider the potential biases in the measurement of happiness, such as social desirability bias or measurement error. Can we discuss these potential biases and how they might affect the findings?

[Amos]: I think were onto something here, Angus. The idea that the flattening pattern is primarily present in the least happy individuals suggests that there may be a threshold effect at play. Perhaps the initial increase in happiness with income is driven by basic needs being met, and once those needs are satisfied, further increases in income don't lead to proportional increases in happiness. This is reminiscent of the concept of diminishing marginal utility, where the marginal utility of additional units of a good decreases as consumption increases. Could we explore this idea further, perhaps using data on basic needs satisfaction, such as access to food and shelter? Additionally, I'm curious about the role of social comparison in shaping emotional well-being. Do people with higher incomes tend to compare themselves to those who are even wealthier, which could lead to decreased happiness even as their absolute wealth increases? We should also consider the potential impact of societal expecta-

tions around income and happiness on our findings. Are there cultural or societal norms that dictate how much happiness one should derive from a given income level? If so, this could influence our results and suggest that the relationship between income and happiness is more complex than we initially thought.

[Angus]: Amos, I'm not convinced that the ceiling effect is the sole explanation for the flattening pattern observed by Kahneman and Deaton. While it's possible that the initial increase in happiness with income is driven by basic needs being met, I think we need to consider other factors, such as the role of expectations and adaptability. For instance, do people with lower incomes have different expectations about their future prospects, which might influence their emotional well-being? Or is it possible that happier individuals are more likely to adapt to changes in their income, leading to continued increases in happiness? I'd like to see more evidence on how expectations and adaptability affect the relationship between income and emotional well-being. Additionally, I think we should also consider the potential biases in the measurement of happiness, such as social desirability bias or measurement error. Can we discuss these potential biases and how they might affect the findings? I'm also concerned that the study may be suffering from the what's the matter with Kansas? problem, where the data is skewed towards the unhappy and less educated, which may not be representative of the general population.

[Amos]: I'm intrigued by the idea that the flattening pattern is primarily present in the least happy individuals. Perhaps this is because people with lower incomes have different reference points for what constitutes happiness, making them more sensitive to small changes in their circumstances. For instance, a person living in poverty might find a small increase in income to be a significant improvement, whereas someone with a higher income might not perceive a similar increase as being as meaningful. This could lead to a non-linear relationship between income and happiness, where small increases in income have a disproportionate impact on happiness for those who are already struggling. Additionally, I think we should also consider the role of social comparison in shaping emotional well-being. Do people with higher incomes tend to compare themselves to those who are even wealthier, which could lead to decreased happiness even as their absolute wealth increases?

[Angus]: I'm concerned that the finding that the flattening pattern is primarily present in the least happy individuals might be due to a selection bias, where those who are already unhappy are more likely to report their unhappiness. Additionally, I think we need to consider the role of social norms and expectations around income and happiness. For instance, in many societies, there is a cultural expectation that income and happiness are strongly correlated, which could lead to a self-reinforcing cycle where people who are unhappy are more likely to report their unhappiness as a result of societal pressure. We should also explore the possibility that the relationship between income and happiness is moderated by factors such as personality traits, such as optimism or pessimism, or by individual differences in values and goals.

[Amos]: Angus, I understand your skepticism, but I think were onto something here. The idea that the flattening pattern is primarily present in the least happy individuals suggests that there may be a threshold effect at play. Perhaps the initial increase in happiness with income is driven by basic needs being met, and once those needs are satisfied, further increases in income don't lead to propor-

tional increases in happiness. This is reminiscent of the concept of diminishing marginal utility, where the marginal utility of additional units of a good decreases as consumption increases. I'd love to explore this idea further, perhaps using data on basic needs satisfaction, such as access to food and shelter. And let's not forget the role of social comparison in shaping emotional well-being. Do people with higher incomes tend to compare themselves to those who are even wealthier, which could lead to decreased happiness even as their absolute wealth increases? We should also consider the potential impact of societal expectations around income and happiness on our findings. Are there cultural or societal norms that dictate how much happiness one should derive from a given income level? If so, this could influence our results and suggest that the relationship between income and happiness is more complex than we initially thought.

[Angus]: I'm not convinced that the ceiling effect is the sole explanation for the flattening pattern observed by Kahneman and Deaton. While it's possible that the initial increase in happiness with income is driven by basic needs being met, I think we need to consider other factors, such as the role of expectations and adaptability. For instance, do people with lower incomes have different expectations about their future prospects, which might influence their emotional well-being? Or is it possible that happier individuals are more likely to adapt to changes in their income, leading to continued increases in happiness? I'd like to see more evidence on how expectations and adaptability affect the relationship between income and emotional well-being. Additionally, I think we should also consider the potential biases in the measurement of happiness, such as social desirability bias or measurement error. Can we discuss these potential biases and how they might affect the findings? I'm also concerned that the study may be suffering from the what's the matter with Kansas? problem, where the data are skewed towards the unhappy and less educated, which may not be representative of the general population.

3 Relevance of Kahneman's last paper to Benefit-Cost analysis

There is a long line of research examining theoretically the relative importance of tax policy and expenditure/regulatory policy in improving social welfare. Mainstream economists have long argued that tax policy is better suited for issues of distribution compared to spending and regulatory programs (Hyland & Zeckhauser, 1979; Kaplow, 2020). Although not disagreeing with the conclusion of a dominant tax policy, Liscow and Sunstein (2024) put forth the practical argument that tax policy is relatively rarely used and that the time may have come to consider the possibility of a two-for-one so to speak. Perhaps the more frequently in-play regulatory policy could be used to generate both desirable efficiency consequences as well as desirable distributional consequences. In an eloquently insightful paper, Viscusi (2024) discusses the issue of what social welfare function to use in any expanded benefit-cost analysis. His discussion emphasizes that even if the social welfare function does not contain direct social welfare enhancing effects of reduced inequality, so that regulatory policy is capable of maximizing potential social welfare, there still is the issue

of what welfare function to use. Viscusi (2024) then proceeds to critique the social welfare function recently suggested that the US Office of Management and Budget use. Our objective in the section to follow is more modest than his. The possibility that regulatory policy can satisfy the minimally desirable outcome of a net social utility gain simply by implicitly transferring income from upper to lower income groups via regulations' effects rests on a necessary amount of diminishing marginal utility. Regulation based implicit progressive income taxation that implicitly redistributes income from upper to lower income groups ideally produce a utility gain for income recipients that exceeds the utility loss of the households whose income is implicitly taxed. We now show that in the simple case that Viscusi (2024) fleshes out the diminishing marginal utility of income needed for social welfare improvement implicitly via regulatory policy is unsupported by the empirical results in Kahneman's last published paper, Killingsworth et al. (2023).

3.1 Empirical results: Data and quantile regression estimates

The starting point in our discussion and interpretation of Kahneman's adversarial collaboration with Killingsworth and Mellers is their realization in Killingsworth et al. (2023) that because of a ceiling effect on the dependent variable the earlier paper of Kahneman and Deaton (2010) was more properly viewed as studying how income affects *un*happiness. Additionally, there was unmodeled heterogeneity that could be addressed in a different regression functional form. This led to the research presented in Killingsworth et al. (2023) that we now discuss and apply to the issue of regulatory reform seeking to rely on possible social utility enhancing effects of income redistribution.

3.1.1 Data used in regression estimates of Killingsworth et al. (2023)

Data used in the focal regression coefficients' estimation of Killingsworth et al. (2023) are publicly available at https://osf.io/qye4a/. They cover 1,725,994 experien ce sampling reports from 33,391 employed US adults current happiness about three times per day over several weeks. Their empirical linear-log regression model admits estimated heterogeneity in the marginal effect of income on happiness by both the level of the dependent variable and the independent variable via a quantile regression with two linear segments joined at (the log of) $100,000 of income.

3.1.2 Quantile regression estimates

The quantile regression approach is a straightforward way to include heterogeneity in estimated marginal effects by the quantiles of the (predicted) dependent variable. (Kniesner et al., 2010; Butler et al., 2023 provide many examples). So as not to overwhelm the reader we will discuss the estimates of the coefficients of the quantile regressions for the 30th and the 70th happiness z-score percentiles. This captures a type of symmetry where we discuss results for persons with estimated happiness that exceeds the lowest 30% and persons with estimated happiness that is less than the highest 30%.

30th Percentile Quantile Regression for Happiness Z-Score Regressed on Log Income

$$Slope\ (t-value)\ below\ Income\ of\ \$100{,}000\ =\ 1.32\ (6.3)\,. \tag{1}$$

$$Slope\ (t-value)\ above\ Income\ of\ \$100{,}000\ =\ 1.21\ (4.3)\,. \tag{2}$$

70th percentile Quantile Regression for Happiness Z-Score Regressed on Log Income

$$Slope\ (t-value)\ below\ Income\ of\ \$100{,}000\ =\ 1.18\ (5.3)\,. \tag{3}$$

$$Slope\ (t-value)\ above\ Income\ of\ \$100{,}000\ =\ 1.92\ (5.5)\,. \tag{4}$$

The two regression results to compare for our topic here are (4) and (1) above, which capture the net welfare effect of moving income from the relatively happy higher-income people to the relatively unhappy lower-income people. The result just described is the opposite of the situation hoped for by proponents of possible social welfare enhancing implicit income transfers inside governmentally funded regulatory policy discussed by Liscow and Sunstein (2024) and is in line with the attendant criticism of implicit regulation based income redistribution by Viscusi (2024).

3.2 Implications for current regulatory policy simply put

Will regulatory policy that implicitly redistributes income from high income relatively happy persons to low income relatively unhappy persons raise total happiness on balance according to the estimates in Kahneman's last published paper, Killingsworth et al. (2023)? No. Or, as quantile regression's best known econometrician, Koenker (2023), wrote concerning the results in Killingsworth et al. (2023): "For the naive utilitarian this (Eqs. (1)–(4)) might suggest redistribution schemes from the sad to the glad, but you would have to be mad to subscribe to that."

4 Conclusion

We have argued here that in Kahneman's co-authored last published paper, Killingsworth et al. (2023), there are unappreciated empirical results relevant to a current question in regulatory policy. Can governmental regulation be not only efficiency enhancing but also simultaneously distributionally welfare enhancing? After explaining the data and empirical well-being hypotheses in Killingsworth et al. (2023) we present an AI chatbot in which their project is discussed and clarified by the late Amos Tversky and Sir Angus Deaton. The final section above fleshes out the empirical implications of quantile regressions in the adversarial collaboration that is the focus of our interest here. The results presented in Killingsworth et al. (2023) support the argument in Viscusi (2024) that distributionally based welfare improvement via regulatory policy is much more nuanced and less obvious than it has been made out to be recently, and our continued belief in TANSTAAFL at the nexus of regulatory policy, income distribution, and social welfare.

Kahneman's final research contribution, represented in the adversarial collaboration with Killingsworth and Mellers (2023), demonstrates both methodical sophistication and empirical precision in resolving apparent contradictions regarding the income-happiness relationship. Their findings reveal systematic heterogeneity in how different population segments experience well-being responses to income—patterns that remained obscured in previous research employing conventional analytical approaches.

The focal collaborative research identified that happiness exhibits differential relationships with income across the well-being distribution: the flattening pattern previously identified by Kahneman and Deaton (2010) appears primarily among the least happy population quintile, while happiness continues increasing with income for the happier majority and even accelerates among those in the highest happiness quantiles. This reconciles previous empirical contradictions while revealing more complex psychological relationships than either original study independently identified.

The empirical findings we discuss here contain significant implications for current debates regarding regulatory policy design and welfare assessment. The evidence that marginal happiness returns to income appear larger among relatively happy high-income individuals than relatively unhappy lower-income individuals challenges conventional diminishing marginal utility assumptions that underlie many redistributive approaches to policy. This suggests that regulatory mechanisms attempting to enhance social welfare through implicit income redistribution may achieve results contradicting their intended purposes.

Our application of computational dialogue techniques to simulate scholarly perspectives from Tversky and Deaton illustrates how empirical findings might be interpreted through complementary psychological and econometric frameworks. The interdisciplinary approach highlights how different analytical traditions illuminate distinctive dimensions of the income-happiness relationship, with psychological perspectives emphasizing cognitive mechanisms and reference dependence, while econometric approaches prioritize measurement validity and statistical relationships.

The methodical evolution demonstrated in Kahneman's final published research project exemplifies several scholarly virtues: willingness to subject previous findings to rigorous reexamination, openness to conceptual innovation through adversarial collaboration, and commitment to empirical precision over theoretical preconceptions. The qualities just mentioned defined Kahneman's transformative contributions to behavioral economics throughout his distinguished career and remain evident in his final published work.

The implications of the findings we focus on here extend beyond academic discourse to practical policy design considerations. The evidence presented in Killingsworth et al. (2023) reinforces the argument that distributional objectives remain most appropriately addressed through tax and transfer mechanisms, while regulatory policy should maintain primary focus on efficiency enhancement. As debates continue regarding appropriate integration of distributional considerations into benefit-cost analysis, well-crafted relevant empirical findings provide essential evidence regarding the complex relationship between economic resources and subjective well-being across different population segments.

In conclusion, Kahneman's final research contribution exemplifies how conceptual innovation and empirical precision can resolve apparent contradictions while generating more sophisticated understanding of complex psychological phenomena. The nuanced patterns revealed through the collaborative research highlight the continuing value of Kahneman's intellectual legacy in advancing our understanding of human judgment, decision-making, and experienced well-being.

5 Appendix

5.1 Deeper examination of the simulated conversation

A central theme in the dialogue is the flattening pattern of happiness observed in the least happy individuals, a finding that reconciles the apparent contradictions between earlier research, including the 2010 collaboration between Kahneman and Deaton, and more recent studies. Tversky approaches this with curiosity, suggesting that psychological mechanisms such as expectations, adaptability, and social comparison may play a role. He raises the possibility that lower-income individuals might have different expectations about their future prospects, which could influence their emotional well-being. Alternatively, he speculates that happier individuals may be more adept at adapting to changes in their income, leading to continued increases in happiness.

Deaton, however, expresses skepticism about the robustness of the findings. He emphasizes the need for more evidence to substantiate claims about expectations and adaptability. He raises concerns about potential biases in measuring happiness, such as social desirability bias and selection bias. The divergence in perspectives underscores the interplay between psychological and economic approaches to understanding subjective well-being. Tversky's focus on theoretical exploration complements Deaton's emphasis on empirical validation, illustrating the value of interdisciplinary dialogue.

Both Tversky and Deaton discuss the concept of diminishing marginal utility, particularly in the context of basic needs satisfaction. Tversky suggests that the initial increase in happiness with income may be driven by fulfilling basic needs, such as access to food and shelter. Once basic needs are satisfied, further increases in income may yield diminishing returns in terms of happiness. This resonates with the economic principle of diminishing marginal utility, where the additional utility derived from each unit of a good decreases as consumption increases.

While Deaton acknowledges the plausibility of Tversky's contention of a basic needs threshold, he calls for more data to substantiate the relationship between basic needs satisfaction and happiness. Deaton also highlights the importance of considering alternative explanations, such as the role of societal norms and cultural expectations in shaping the relationship between income and well-being. Their intellectual exchange underscores the necessity for a comprehensive approach integrating psychological and economic perspectives to understand the factors influencing subjective well-being.

Tversky introduces the idea that social comparison may play a significant role in shaping emotional well-being. He posits that higher-income individuals may compare themselves to even wealthier peers, potentially leading to decreased happiness despite absolute income gains. The Tversky perspective highlights the importance of relative income and societal expectations in understanding subjective well-being.

Building on the social comparisons idea, Deaton suggests that cultural and societal norms around income and happiness could influence the findings. For instance, societal pressures and expectations may create a self-reinforcing cycle where individuals report unhappiness due to perceived discrepancies between their income and societal benchmarks. The discussion highlights the complexity of the relationship between income and well-being, suggesting that it is shaped not only by individual factors but also by broader social and cultural dynamics.

5.2 Concerns over method and biases

A key focus of Deaton's contributions to the dialogue is the method's rigor in the 2023 study. He raises concerns about potential biases in measuring happiness, such as social desirability bias, measurement error, and the "What's the Matter with Kansas?" problem, where data may be skewed toward less happy or less educated individuals. Deaton's concerns likely stem from broader commitment to ensuring the validity and reliability of empirical findings, a hallmark of his work in economics.

Although less focused on method critiques, Tversky emphasizes the need to explore alternative explanations and psychological mechanisms. The difference in emphasis illustrates the complementary nature of their perspectives, with Tversky prioritizing theoretical exploration and Deaton emphasizing empirical validation. Together, their two approaches provide a more comprehensive understanding of the relationship between income and subjective well-being.

5.3 Interdisciplinary insights and constructive debate

The dialogue between the two bots exemplifies the value of interdisciplinary exchange in addressing complex questions about income and well-being. Tversky's psychological insights into heuristics, expectations, and social comparison complement Deaton's economic focus on measurement, data quality, and policy implications. Together, their bots' exchanges highlight the multifaceted nature of subjective well-being and the importance of diverse approaches to its study.

Despite their differing perspectives, the dialogue is characterized by a shared curiosity and commitment to understanding the complex interaction between income-happiness relationship. The exchange is marked by constructive debate, with each participant building on the other's ideas and raising thoughtful questions. Their collaborative spirit reflects their intellectual rigor and openness that define their respective fields.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s11166-025-09459-5.

Funding Open access funding provided by SCELC, Statewide California Electronic Library Consortium

Declarations

Competing interests None

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Adler, M. (2016). Behavioral economics, happiness surveys, and public policy. *Journal of Benefit-Cost Analysis*, *7*(1), 196–219.

Anthis, J. R., Liu, R., Richardson, S. M., Kozlowski, A. C., Koch, B., Brynjolfsson, E., Evans, J., & Bernstein, M. S. (2025). LLM social simulations are a promising research method. *Computer Science Human-Computer Interaction*.

Benjamin, D. J., Heffetz, O., Kimball, M. S., & Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose? *American Economic Review, 102*(5), 2083–2110.

Bryson, A., & MacKerron, G. (2016). Are you happy while you work? *The Economic Journal*, *127*(599), 106–125.

Butler, R., Butler, M., & Wilson, B. (2023). Advanced statistics for health research. *Journal of Benefit-Cost Analysis, 15*(2), 224–251.

Clark, A., Frijters, P., & Shields, M. (2008). Relative income, happiness, And utility: An explanation for the Easterlin paradox And other puzzles. *Journal of Economic Literature*, *46*(1), 95–144.

Gonzalez-Bonorino, A., Capra, M., & Pantoja, E. (2025). LLMs model Non-WEIRD populations: Experiments with synthetic cultural agents.

Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from homo silicus?.

Hyland, A., & Zeckhauser, R. (1979). Distributional objectives should affect taxes but not program choice or design. *Scandinavian Journal of Economics*, *81*(2), 264–284.

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.

Kahneman, D., & Deaton, A. (2010). High income improves evaluation of life but not emotional well-being. *Proceedings of the National Academy of Sciences*, *107*(38), 16489–16493.

Kahneman, D., & Krueger, A. (2006). Developments in the measurement of subjective well-being. *Journal of Economic Perspectives*, *20*(1), 3–24.

Kahneman, D., & Sugden, R. (2005). Experienced utility as a standard of policy evaluation. *Environmental and Resource Economics*, *32*(1), 161–181.

Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *The Quarterly Journal of Economics*, *112*(2), 375–405.

Kahneman, D., Krueger, A., Schkade, D., Schwarz, R., & Stone, A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, *306*(5702), 1776–1780.

Kaplow, L. (2020). A unified perspective on efficiency, redistribution, and public policy. *National Tax Journal*, *73*(2), 429–472.

Killingsworth, M. A. (2021). Experienced well-being rises with income, even above $75,000 per year. *Proceedings of the National Academy of Sciences of the United States of America, 118*(4):e2016976118.

Killingsworth, M., Kahneman, D., & Mellers, B. (2023). Income and emotional well-being: A conflict resolved. *Proceedings of the National Academy of Sciences*, *120*(10) e2208661120.

Koenker, R. (2023). How quantile regression rescued happiness—for the rich: An R vinaigrette.

Kniesner, T. J., Viscusi, W. K., & Ziliak, J. P. (2010). Policy relevant heterogeneity in the value of a statistical life: New evidence from panel data quantile regressions. *Journal of Risk and Uncertainty*, *40*(1), 15–31.

Krueger, A. B., Kahneman, D., Schkade, D., Schwarz, N., & Stone, A. A. (2009). National time accounting: The currency of life. In A. B. Krueger (Ed.), *Measuring the subjective Well-Being of nations: National accounts of time use and Well-Being* (pp. 9–86). University of Chicago Press.

Liscow, Z., & Sunstein, C. R. (2024). Efficiency vs. welfare in benefit–cost analysis: The case of government funding. *Journal of Benefit-Cost Analysis, 15*(2), 224–251.

Neve, J., Ward, G., Keulenaer, F., Landeghem, B., Kavetsos, G., & Norton, M. (2018). The asymmetric experience of positive and negative economic growth: Global evidenceusing subjective well-being data. *The Review of Economics and Statistics, 100*(2), 362–375.

Pérez-Truglia, R. (2020). The effects of income transparency on well-being: Evidence from a natural experiment. *American Economic Review*, *110*(4), 1019–1054.

Steptoe, A., Deaton, A., & Stone, A. (2015). Subjective wellbeing, health, and ageing. *The Lancet*, *385*(9968), 640–648.

Stevenson, B., & Wolfers, J. (2013). Subjective well-being and income: Is there any evidence of satiation? *American Economic Review*, *103*(3), 598–604.

Ubel, P., Loewenstein, G., & Jepson, C. (2005). Disability and sunshine: Can hedonic predictions be improved by drawing attention to focusing illusions or emotional adaptation? *Journal of Experimental Psychology Applied*, *11*(2), 111–123.

Viscusi, W. (2024). Why office of management and budget's (OMB) social welfare function is not society's social welfare function. *Journal of Benefit Cost Analysis*, *15*(2), 252–275.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.