

INTERNATIONAL
CENTER FOR
PUBLIC POLICY

—
WITHIN THE
PUBLIC FINANCE
RESEARCH CLUSTER

Known Unknowns: Knowledge, Confidence, and Public Opinion on Water Reuse

Alberto Chong
Marco D'Errico

October 2026



ANDREW YOUNG SCHOOL
OF POLICY STUDIES

Tel: 404-413-0233

Website: pfrc.gsu.edu

Address:
55 Park Place NE
7th Floor
Atlanta, GA 30303

Mail:
Public Finance Research Cluster
P.O. Box 3992
Atlanta, GA 30302-3992

International Center for Public Policy Andrew Young School of Policy Studies

The Andrew Young School of Policy Studies was established at Georgia State University with the objective of promoting excellence in the design, implementation, and evaluation of public policy. In addition to four academic departments, including economics and public administration, the Andrew Young School houses nine leading research centers, including the International Center for Public Policy.

The mission of the International Center for Public Policy (ICePP) is to provide academic and professional training, applied research, and technical assistance in support of sound public policy and sustainable economic growth in developing and transitional economies.

ICePP is recognized worldwide for its efforts in support of economic and public policy reforms through technical assistance and training around the world. This reputation has been built serving a diverse client base, including the World Bank, the U.S. Agency for International Development (USAID), the United Nations Development Programme (UNDP), finance ministries, government organizations, legislative bodies, and private sector institutions.

The success of ICePP reflects the breadth and depth of its in-house technical expertise. The Andrew Young School's faculty are leading experts in economics and public policy and have authored books, published in major academic and technical journals, and have extensive experience in designing and implementing technical assistance and training programs. Andrew Young School faculty have been active in policy reform in over 70 countries around the world. Our technical assistance strategy is not merely to provide technical prescriptions for policy reform, but to engage in a collaborative effort with host governments and donor agencies to identify and analyze the issues at hand, arrive at policy solutions, and implement reforms.

ICePP specializes in four broad policy areas:

- Fiscal policy (e.g., tax reforms, public expenditure reviews)
- Fiscal decentralization (e.g., reform, intergovernmental transfer systems, urban finance)
- Budgeting and fiscal management (e.g., local, performance-based, capital, and multi-year budgeting)
- Economic analysis and revenue forecasting (e.g., micro-simulation, time series forecasting)

For more information about our technical assistance activities and training programs, please visit our website at icepp.gsu.edu or contact us at paulbenson@gsu.edu.

Known Unknowns: Knowledge, Confidence, and Public Opinion on Water Reuse

Alberto Chong and Marco D'Errico*

October 2026

Abstract

Public acceptance can constrain treated-wastewater reuse, but less is known about whether citizens provide a substantive view at all. An influential account of opposition to contested technologies predicts that the strongest opponents know the least while believing they know the most. We test that joint prediction using a knowledge–attitudes–practices survey of 1,069 households in Barbados. The survey separately measures demonstrated topic knowledge, self-assessed knowledge, willingness to use reclaimed water, and explicit ‘don't know’ responses. Among respondents who provide a complete willingness profile, demonstrated knowledge is essentially unrelated to acceptance, whereas self-assessed knowledge is positively associated with acceptance. The predicted knowledge-overconfidence configuration therefore does not appear in this setting. The strongest knowledge gradient instead lies on the response margin: a one-standard-deviation increase in demonstrated knowledge is associated with a 5.8-percentage-point lower probability of giving at least one ‘don't know’ response, against a base rate of 22.6 per cent. The relationship reappears in three sectoral attitude batteries and attenuates by about 19 per cent after controlling for out-of-module nonresponse. The evidence therefore points away from confidently mistaken opposition and toward a sizeable group that is less knowledgeable and more likely to express uncertainty.

Keywords: water reuse, public acceptance, stated preferences, surveys, small island developing states

JEL Codes: Q25, Q53, D83, D91, O13

* Chong (achong6@gsu.edu): Department of Economics, Georgia State University, Atlanta GA, USA and Department of Economics, Universidad del Pacifico, Lima, Peru; D'Errico (mderrico@gcfund.org) Independent Evaluation Office, Green Climate Fund, Songdo Business District, South Korea.

We are grateful to Michell Dong, Susumu Yoshida, and Luisa Zanforlin for comments and suggestions and to Young-Yun Choi for excellent research assistance. All remaining errors are ours.

1. Introduction

Water-scarce societies increasingly consider treated-wastewater reuse as part of their climate-adaptation portfolios. Treatment technologies are mature, but technical feasibility does not guarantee implementation: public concern about health risks, contamination, and institutional trust can constrain otherwise viable reuse schemes. Understanding not only whether citizens support reuse, but also whether they provide a substantive answer at all, is therefore relevant to program design. One influential account of opposition to contested technologies focuses on knowledge miscalibration. Fernbach et al. (2019) show for genetically modified foods that, as opposition becomes more extreme, objective knowledge declines while perceived understanding rises; Light et al. (2022) find related patterns across several controversial scientific issues, while also documenting important exceptions. The empirical claim is therefore issue specific. Its policy relevance is potentially important: if opposition is coupled with unrecognized ignorance, simply supplying more facts may have limited traction because the intended audience does not perceive an information deficit.

Water reuse offers a useful boundary test because low observed support can arise on two distinct margins. At the intensive margin, citizens may hold a substantive view and oppose reuse. At the response margin, they may be unwilling or unable to state a directional view. These states are substantively different. Strong, confidently held opposition may be difficult to move; an explicit ‘don’t know’ response may instead reflect recognized uncertainty, limited information, reluctance to guess, satisficing, or item-specific difficulty. The appropriate communication strategy depends on which margin is empirically important. The existing water-reuse literature provides rich evidence on the direction and intensity of attitudes. Reviews emphasize disgust, perceived health risk, trust, communication, and end use, while showing that their relative importance varies across places and applications (Nkhoma et al., 2021; Hurlimann and Dolnicar, 2016). Recent systematic evidence likewise emphasizes public education, drought conditions, institutional arrangements, and other context-specific factors in successful reuse implementation (Sardana et al., 2025). Recent survey evidence from small U.S. communities finds that knowledge, perceived risks and benefits, trust, and social norms are associated with acceptance, and that prior exposure can operate through perceived risk (Wang, et al., 2026). These studies reinforce the importance of knowledge for the intensive margin but do not make the response margin a central outcome. A parallel valuation literature estimates willingness to pay or use and studies how information changes stated

acceptance. Less attention has been paid to whether a respondent can state a preference in the first place. This omission matters because survey-methodology and contingent-valuation research has long shown that ‘don’t know’ responses are not generically interchangeable with missing observations or substantive yes/no answers (Haener and Adamowicz, 1998; Wang, 1997; Elkjær and Wlezien, 2025).

We examine both margins using a knowledge–attitudes–practices survey of 1,069 households across 30 enumeration districts in Barbados, a highly water-scarce small island developing state. Caribbean evidence on public attitudes toward wastewater reuse remains thin: Roopnarine et al. (2023) provide an important Trinidad and Tobago precedent. The Barbados survey adds a household-scale sample and, crucially for our question, measures demonstrated topic knowledge and self-assessed knowledge separately while allowing respondents to answer 'don't know' on four willingness items and on three additional sectoral batteries. The first result is a boundary test for the knowledge-overconfidence account. Objective and self-assessed knowledge are only moderately aligned ($\rho = 0.32$), so the data contain room for miscalibration. Yet among respondents who provide a complete willingness profile, objective knowledge is essentially flat across the opposition–acceptance spectrum. In the fully controlled specification, its coefficient is -0.022 standard deviations ($p = 0.55$), with a 95 per cent confidence interval of $[-0.094, 0.049]$. Self-assessed knowledge, by contrast, is positively associated with acceptance ($+0.091$, $p = 0.029$). The Fernbach-style interaction between objective knowledge and opposition is also near zero, and the standardized difference between self-assessed and objective knowledge tilts weakly in the opposite direction from the hypothesis.

The second result shifts the focus from the direction of opinion to whether a respondent supplies a complete willingness profile. Twenty-two point six per cent of respondents answer 'don't know' to at least one willingness item. A one-standard-deviation increase in objective knowledge is associated with a 5.8-percentage-point lower probability of this response pattern after controls and parish fixed effects; the wild-cluster bootstrap p-value is 0.008. Self-assessed knowledge independently predicts lower abstention as well. Respondents who do not provide a complete willingness profile therefore tend, on average, both to demonstrate less topic knowledge and to describe themselves as less informed.

The response-margin association is not confined to one battery. Demonstrated knowledge predicts 'don't know' responses on perceived challenges in agriculture, tourism and business, with coefficients of -0.143 , -0.132 and -0.166 respectively. The willingness abstainers overlap only weakly with abstainers in these sectoral batteries, although the three sectoral batteries are themselves strongly correlated. We therefore interpret the exercise as replication across two substantively distinct types of judgment, personal willingness and third-party sectoral assessment, not as four independent replications.¹

Two threats to interpretation are especially important. First, the knowledge index may partly capture general response capacity or survey fluency. An out-of-module nonresponse measure is related to demonstrated knowledge; adding it to the main specification attenuates the coefficient from -0.058 to -0.047 , or about 19 per cent, but does not eliminate the association. We do not interpret that attenuation as a causal decomposition, and the auxiliary measure is not a pure measure of response capacity because it includes an income item that may also reflect privacy or sensitivity. Second, because the direction regressions condition on providing a complete willingness profile, they are selected by construction. Two deliberately polar common-value imputations—assigning every abstainer either the maximum or the minimum acceptance score—move the demonstrated-knowledge coefficient to -0.084 and $+0.118$, respectively. These are sensitivity checks, not formal missing-data bounds, because sharper worst-case assignments could vary across abstainers with knowledge and covariates.

These cross-sectional data do not show that an information campaign would create opinions or increase acceptance. They do identify a sizeable group whose explicit uncertainty coexists with lower demonstrated and self-assessed knowledge, and they generate a testable prediction: an effective information intervention may first reduce the response margin of 'don't know' answers, while its effect on the direction of newly expressed opinions is theoretically ambiguous.²

The paper makes three contributions. First, it provides a boundary test of the knowledge-overconfidence account by examining its joint objective–subjective knowledge signature in a water-reuse setting; the evidence does not reproduce the pattern documented for genetically

¹ The association is also driven primarily by recognition/content-production items and is concentrated among respondents without tertiary education and those aged sixty and over.

² A randomized information experiment or panel survey that measures both margins would be required to test that prediction.

modified foods and several other contested issues. Second, it treats stated acceptance as a two-part outcome, that is, whether a respondent supplies a substantive response and, conditional on doing so, the direction of that response, including connecting water-reuse research to work on non-attitudes, response uncertainty and contingent valuation. Third, it shows empirically that the extensive margin contains a strong knowledge gradient that would be obscured if explicit 'don't know' responses were simply deleted. The result is not that abstention is a pure measure of latent non-opinion, but that it is an informative response state deserving separate analysis.

2. Related Literature

This paper lies at the intersection of several strands of the literature: public acceptance of water reuse, information and stated valuation, knowledge calibration, and survey research on 'don't know' responses and non-attitudes. The central synthesis is not that abstention reveals a latent non-attitude, but that public-opinion measurement has two observable margins, whether a respondent supplies a substantive response and, conditional on doing so, what that response is.

Research on water reuse has identified a recurring set of barriers without reducing acceptance to a single universal mechanism. Reviews emphasize disgust and contamination intuitions, perceived health and quality risks, trust in the institutions that design and operate reuse systems, and the credibility and framing of communication (Po et al., 2003; Russell and Lux, 2009; Rozin et al., 2015; Ross et al., 2014; Nkhoma et al., 2021). Comparative evidence also shows substantial variation by location, water source and end use. For instance, respondents are typically more accepting of uses involving less bodily contact, and the ranking of alternative sources changes across countries and applications (Hurlimann and Dolnicar, 2016). Recent evidence reinforces this contextuality. Sardana et al., (2025) synthesize 232 peer-reviewed studies of water-reuse implementation and identify public education, drought conditions, regulation and institutional capacity among recurring implementation factors; Wang et al., (2026) find that knowledge, perceived risks and benefits, trust and social norms are associated with acceptance in small U.S. communities. The relevant lesson for our design is that knowledge is one component of a multidimensional acceptance process, not a universal determinant of support.

A complementary literature argues that acceptance is better understood as a social process than as a one-time attitude snapshot. Al-Saidi (2021) organizes that process around public perception, politicisation, individual acceptance and adaptation in use, stressing interactions among source,

technology and end use. Ormerod and Scott (2013), for example, show that willingness to drink reclaimed water is closely related to trust in the professional and political institutions governing supply. This process perspective makes the response margin conceptually relevant: citizens can be at different stages of engagement with a technology, and a respondent may be willing to acknowledge uncertainty. At the same time, the existing literature does not imply that 'don't know' is equivalent to non-opinion; the survey-methodology literature reviewed below provides the necessary qualification.

The evidence base is also geographically uneven. Much of the canonical literature comes from Australia, the United States and Europe, while Caribbean evidence remains extremely limited. Roopnarine et al. (2023) study 146 stakeholders in Trinidad and Tobago in a KAP and willingness-to-pay assessment and document the coexistence of stated knowledge, health concerns and uneven attitudes toward wastewater reuse. Our study therefore does not claim a first Caribbean analysis. Its contribution is to bring a household-scale Barbados sample and, more importantly, a separate analysis of demonstrated knowledge, perceived knowledge and explicit response uncertainty.

An economic strand values reclaimed water and studies whether information shifts stated demand. Stated-preference studies report positive willingness to pay or use under specified safety and price conditions and often find that information is associated with greater acceptance (Tsagarakis and Georgantzis, 2003; Menegaki, et al., 2007; Dolnicar, et al., 2011). Experimental evidence is more informative about mechanism. Fielding and Roiko (2014) find that brief information about potable recycled water generally increases comfort and support, while some pollutant-focused information produces more ambivalent responses. More recent work similarly shows that educational content can increase knowledge and acceptance, while reviews caution against treating an information-deficit model as a complete account of public acceptance (Moesker, et al., 2024). The implication for this paper is not that information must raise support, but that information may affect both the ability to answer and the direction of an answer. Most willingness-to-pay and willingness-to-use analyses condition on respondents providing a substantive valuation. Environmental economics has long recognized that this conditioning can be consequential. Haener and Adamowicz (1998) find that contingent-valuation 'don't know' respondents differ systematically from both yes and no respondents, while Wang (1997) develops a random-valuation model that treats 'don't know' as preference uncertainty instead of automatically recoding or deleting it. Recent experimental evidence also shows that offering an explicit 'don't know' option can change measured policy

preferences, especially among lower-information respondents and on lower-salience issues (Elkjær and Wlezien, 2025). These literatures supply an economic and measurement precedent for our response-margin analysis: whether a stated preference is available for valuation is itself an empirical outcome.

Our main directional hypothesis comes from work on the calibration of demonstrated and subjective knowledge. Kruger and Dunning (1999) provide the classic evidence that poor performers can have difficulty recognizing their own errors, but the empirical prediction tested here is narrower than a generic 'Dunning–Kruger effect.' Fernbach et al. (2019) identify a specific joint configuration for genetically modified foods: as opposition becomes more extreme, demonstrated knowledge declines while perceived understanding rises. Light et al. (2022) extend the analysis across several controversial scientific issues. The theoretically relevant object is therefore the covariance structure among demonstrated knowledge, perceived knowledge and opposition, not simply a large value of a difference score. Our direct interaction test follows this logic; the difference score is reported only as a descriptive summary. The pattern is not universal. Fernbach et al. (2019) themselves report that the configuration did not emerge for climate-change attitudes, and Light et al. (2022) document issue-level variation. Water reuse is therefore a useful boundary case: its opposition can involve contamination and disgust, but the issue may be less identity-laden and less politically entrenched than some of the controversies in the overconfidence literature. We reproduce the joint tests as directly as the survey permits and interpret any divergence as evidence about the scope of the empirical regularity, not as evidence that the underlying psychological mechanism has been disproved.

The survey-methodology literature provides the final component of our framework. Converse (1976) documents the prevalence and predictors of 'no opinion' responses in public-opinion surveys, while Krosnick (1991) develops the satisficing account. Krosnick et al. (2002) show experimentally that explicitly offering a no-opinion option can attract respondents with lower cognitive skills or effort, so an explicit 'don't know' response cannot automatically be equated with a latent non-attitude. Luskin and Bullock (2011), however, show that political knowledge 'don't know' responses also contain information about genuine knowledge deficits. Taken together, this literature motivates treating explicit abstention as an informative but behaviorally mixed response state: it can reflect uncertainty, low information, satisficing, reluctance to guess, or a response strategy. Recent measurement research reinforces both sides of that qualification.

Kuha et al. (2018) find experimentally that probing initial ‘don’t know’ answers reduces their frequency but can worsen measurement quality for attitude items. Elkjær and Wlezien (2025) show that offering a ‘don’t know’ option can materially affect estimated policy preferences, particularly among lower-information respondents and on low-salience issues. We therefore use ‘abstention’ only as an operational label for an explicit ‘don’t know’ response, test whether it is systematically related to knowledge, and examine whether the relationship survives an out-of-module nonresponse control. The contribution is not to identify the latent meaning of each ‘don’t know’ response, but to show that deleting this response category would remove a systematic knowledge gradient from the analysis.

The literature yields two distinct empirical questions. First, if the knowledge-overconfidence account generalizes to water reuse, greater opposition should be associated with lower demonstrated knowledge and higher perceived knowledge, with the relationship between the two becoming negative among stronger opponents. This is the intensive-margin prediction. Second, the survey permits a separate question that the overconfidence account does not directly address, in particular, whether knowledge predicts the probability that respondents provide a substantive answer at all. We therefore treat the response margin as an empirical outcome in its own right. A negative knowledge–abstention association would be consistent with lower-information respondents being more likely to express uncertainty, but by itself would not establish why they abstain or imply that information would move their eventual opinions toward support.

3. Setting and Data

Barbados is among the most water-scarce countries in the world on a per-capita basis, with roughly 280 cubic meters of renewable internal freshwater available per person per year (World Bank, 2023), well below the conventional threshold of 500 cubic meters for absolute water scarcity. The island has no significant perennial surface water; public supply depends almost entirely on groundwater abstracted from a coral limestone aquifer, exposing the resource to over-abstraction and, near the coast, to saline intrusion. Observed climate records show an increasing incidence of dry periods, and projections point to declining annual precipitation and a lengthening dry season. Sanitation coverage is limited: two municipal treatment plants serve on the order of three per cent of the population, and the remainder of the island relies on septic systems that, where poorly maintained, threaten the same aquifer on which potable supply depends. Treating wastewater to

tertiary standards could substitute for groundwater in non-potable uses, principally agricultural irrigation, and could support aquifer recharge. The survey was designed to measure public attitudes toward that potential, but the analysis below is not a causal evaluation of any reuse intervention.

The data come from a knowledge–attitudes–practices household survey administered face-to-face by trained enumerators using computer-assisted personal interviewing, which enforced the instrument’s skip logic. One adult respondent answered on behalf of each household. The instrument covers household water practices, knowledge of water and reuse, attitudes toward reclaimed water and sector policy, perceptions of reuse in agriculture, tourism and business, and demographics. The realized sample is 1,069 households across 30 enumeration districts in four parishes. The enumeration district is the smallest clean geographic unit in the data and the level at which we cluster all standard errors; because thirty clusters is few for conventional asymptotic inference, we accompany headline coefficients with a wild-cluster bootstrap p-value using Rademacher weights and 1,999 replications (Cameron and Miller, 2015; Roodman, et al., 2019). The survey was fielded as the baseline of a program evaluation whose two community arms were formed purposively, not by randomization; we therefore pool the arms throughout, include an arm indicator and parish fixed effects among the controls, and draw no inference from the arm contrast. Our estimates are conditional associations, not treatment effects, and we use no causal language for them.

4. Measuring Knowledge, Confidence, and Opinion

The empirical exercise keeps four constructs separate: demonstrated topic knowledge, perceived knowledge, the direction of stated willingness, and the explicit decision to answer ‘don’t know.’ This distinction matters because the knowledge-overconfidence hypothesis concerns the joint relation between demonstrated and subjective knowledge and opposition, whereas the response-margin analysis concerns the fourth construct. Throughout the paper, ‘opinion-holder’ and ‘abstainer’ are convenient operational labels for a complete substantive response and at least one explicit ‘don’t know’ response; they are not claims that latent attitudes literally do or do not exist.

The demonstrated-knowledge index counts how many of ten items elicit a prespecified substantive answer. Four concern institutions and instruments: whether the respondent can describe water reuse, is aware of the planned treatment-plant upgrade, is aware of national water policies, and is aware of reuse legislation. Six are recognition or production items requiring substantive content:

naming a commercial, agricultural, and residential benefit of reuse; a waterborne illness; an environmental consequence of poor water management; and a problem in national water management. Multiple-response items are scored one when any prespecified substantive option is named, excluding ‘none’ and ‘don’t know.’ We refer to the resulting measure as demonstrated topic knowledge and not a conventional factual-knowledge test. It averages 3.47 of 10 (s.d. 2.22) and has Cronbach’s $\alpha = 0.70$. The recognition items are on average easier (mean difficulty 0.50 versus 0.11) and more discriminating (mean item-rest correlation 0.44 versus 0.22) than the institutional-awareness items, although the pattern is not uniform. Because open-ended recognition can also load on verbal fluency and response capacity. Please, see Table 1.

Subjective knowledge is the mean of two items asking respondents to rate how much they know about current and possible water-reuse practices, reverse-coded so that higher values denote greater self-rated knowledge (mean 1.58 on a 1–4 scale, $\alpha = 0.93$, $n = 1,029$). For comparability with the overconfidence literature, we report a miscalibration index equal to standardized subjective knowledge minus standardized demonstrated knowledge. Demonstrated and subjective knowledge are positively but weakly correlated ($\rho = 0.32$), and 41.9 per cent of respondents lie above zero on the difference score. Because difference scores can inherit variation from both components, the primary tests use the two components separately and reproduce the interaction structure underlying Fernbach et al. (2019); the difference score is descriptive.

Table 1. Objective-Knowledge Items

Item	Type	Share correct	Item-rest corr.
Can describe what water reuse is	Factual	0.198	0.357
Aware of treatment-plant upgrade	Factual	0.189	0.195
Aware of national water policies	Factual	0.041	0.171
Aware of reuse legislation	Factual	0.029	0.161
Names a commercial benefit of reuse	Recognition	0.395	0.643
Names an agricultural benefit of reuse	Recognition	0.499	0.622
Names a residential benefit of reuse	Recognition	0.475	0.561
Names a waterborne illness	Recognition	0.526	0.313
Names an environmental consequence	Recognition	0.486	0.151
Names a water-management problem	Recognition	0.630	0.335

Notes: $N = 1,069$. Index mean 3.47, s.d. 2.22, Cronbach’s $\alpha = 0.70$. Item-rest correlation is the correlation between the item and the sum of the remaining nine.

Acceptance is the mean of four items asking whether the respondent would use reclaimed water if it were reliable, safe, inexpensive, and delivered through a separate pipe. Each uses a four-point response scale reverse-coded so that higher values denote greater willingness (mean 3.06, $\alpha = 0.91$). Each item also permits an explicit 'don't know' response. The acceptance index therefore describes the direction of stated willingness; it is not defined for respondents who do not provide a complete profile in the main direction analysis. The response-margin outcome is an indicator equal to one if the respondent answers 'don't know' on at least one of the four willingness items. It equals one for 22.6 per cent of the sample. A stricter definition, requiring 'don't know' on all four items, applies to 7.3 per cent and is reported as a robustness check. A respondent who answers 'don't know' on some but not all items can still receive a descriptive acceptance index, so that index is observed for 990 households. The direction regressions, however, use only respondents recording no 'don't know' on any of the four items, ensuring a complete acceptance measure for every observation in that analysis. There are 827 such respondents, of whom 826 enter the direction regressions because one otherwise complete respondent is missing the acceptance items for another reason. Coding that case as an abstainer changes the response rate only trivially and does not affect the results.

Three further batteries permit the same explicit response and are used below: perceived challenges of reclaimed water use in agriculture (37.6 per cent abstain), tourism (58.4 per cent) and business (57.1 per cent). These batteries let us ask whether the knowledge-abstention relationship is peculiar to personal willingness or extends to judgments about sectoral consequences. Table 2 reports summary statistics for all constructs and covariates.

Table 2. Summary Statistics

Variable	Mean	S.D.	Min	Max
Demonstrated knowledge (0–10)	3.47	2.22	0	10
Subjective knowledge (1–4)	1.58	0.74	1	4
Overconfidence ($z(\text{SK}) - z(\text{OK})$)	−0.02	1.17	−3.72	4.84
Acceptance index (1–4)	3.06	0.87	1	4
Abstains, willingness battery (0/1)	0.226	0.419	0	1
Abstains on all four items (0/1)	0.073	0.260	0	1
Abstains, agriculture battery (0/1)	0.376	0.485	0	1
Abstains, tourism battery (0/1)	0.584	0.493	0	1
Abstains, business battery (0/1)	0.571	0.495	0	1

Female respondent	0.528	0.499	0	1
Age (years)	55.1	17.3	18	96
Tertiary education	0.329	0.470	0	1
Any farming (incl. kitchen garden)	0.140	0.347	0	1
Water interruption, past month	0.075	0.263	0	1

Notes: N = 1,069 households in 30 enumeration districts; subjective knowledge is observed for 1,029 and the acceptance index for 990.

5. Empirical Strategy

Our analysis is descriptive by design. The estimates are conditional associations between demonstrated knowledge, perceived knowledge, and two observable response margins; none of the regressions is interpreted causally. We organize the analysis as a two-part description of stated opinion. The first part asks whether a respondent provides a complete substantive willingness profile. The second asks, conditional on doing so, how favorable that profile is. This ordering prevents the direction of opinion from being conflated with the existence of a stated response, while making explicit that the second part is a selected subsample. For the direction of opinion, we restrict attention to complete opinion-holders and regress standardized acceptance on standardized demonstrated knowledge and standardized subjective knowledge, first separately and then jointly, with and without demographic controls and parish fixed effects:

$$Acceptance(i) = \alpha + \beta \cdot OK(i) + \gamma \cdot SK(i) + X(i)' \delta + \mu(p) + \varepsilon(i) \quad (1)$$

The knowledge-overconfidence account is not a prediction about the coefficient on demonstrated knowledge alone. It predicts a joint configuration in which demonstrated knowledge falls and perceived knowledge rises as opposition becomes more extreme, with the relation between the two becoming negative at high levels of opposition. We therefore supplement the direction regression with three direct tests: the Fernbach-style interaction between demonstrated knowledge and opposition in a regression for subjective knowledge; each knowledge component separately on opposition; and, as a descriptive summary only, the standardized subjective-minus-demonstrated difference on opposition. This keeps the primary inference tied to the joint configuration.

For the response margin, we estimate linear probability models of explicit ‘don’t know’ abstention on the same knowledge measures over the full sample and report logit average marginal effects as a functional-form check. Controls are sex, age, tertiary education, farming status, recent supply interruption, and the evaluation-arm indicator; parish fixed effects absorb broad geographic


differences. Standard errors are clustered on the 30 enumeration districts. Because thirty clusters is modest for asymptotic inference, headline coefficients are accompanied by wild-cluster bootstrap p-values using Rademacher weights and 1,999 replications.

Household survey weights, alternative index constructions, and a stricter abstention definition are reported as robustness checks. Conditioning the direction regression on a complete substantive response creates a transparent selection issue: if knowledge predicts the probability of abstention, the observed opinion-holder sample is selected on a variable related to the regressor of interest. We do not fit a parametric selection model because the survey provides no credible exclusion restriction. Instead, Section 7 reports two deliberately polar common-value imputations that assign every abstainer the minimum or maximum acceptance value permitted by the scale. These stress tests illustrate sensitivity under simple extreme scenarios, but they are not formal partial-identification bounds because the coefficient-extremizing assignment could vary across abstainers with knowledge and covariates.

6. Results

Objective and subjective knowledge are positively but weakly related. Figure 1 plots mean self-rated knowledge across quintiles of the demonstrated-knowledge index: self-assessment rises from 1.38 in the lowest quintile to 1.96 in the highest, while demonstrated knowledge rises from 0.58 to 6.70 substantive items over the same range. The correlation of 0.32 shows meaningful but far from complete alignment. This is a necessary precondition for the overconfidence hypothesis, not evidence for it: the hypothesis additionally requires miscalibration to be concentrated among opponents. Importantly, respondents with low demonstrated knowledge do not generally describe themselves as highly informed; the sample mean for subjective knowledge is only 1.58 on a four-point scale, and the modal response is at the low end.

Figure 1. Calibration of Self-rated against Demonstrated Knowledge



Points are quintile means of self-rated knowledge (left axis, circles) and of demonstrated knowledge (right axis, squares) across quintiles of the objective index; whiskers are 95 per cent confidence intervals. Self-assessment rises only mildly across a wide range of demonstrated knowledge.

Table 3 reports the direction regressions among complete opinion-holders. Demonstrated knowledge is close to zero in every specification: -0.022 bivariate ($p = 0.65$), -0.058 when subjective knowledge is added without controls ($p = 0.22$), and -0.022 with demographic controls and parish fixed effects ($p = 0.55$). The 95 per cent confidence interval in the fully controlled specification is $[-0.094, 0.049]$ standard deviations. Thus, the data provide no evidence of the monotone negative knowledge–acceptance relation required by the confidently mistaken-opponent account, and the preferred estimate is centered close to zero. This is a conditional result among respondents who supplied a complete willingness profile; it does not identify the population relationship that would obtain if abstainers had observed latent acceptance values. Subjective knowledge is associated with the direction of opinion, but in the opposite direction from the hypothesis: the fully controlled coefficient is $+0.091$ standard deviations ($p = 0.029$). Respondents who rate themselves as more informed are, on average, more accepting, conditional on demonstrated knowledge and the included covariates. Because the design is cross-sectional, this coefficient should not be read as the effect of confidence on acceptance; its relevance is that the sign does not match the proposed overconfidence configuration.

Table 3. Knowledge and the Direction of Opinion (Opinion-Holders)

Dependent variable: acceptance (z)	(1) OK only	(2) SK only	(3) OK + SK	(4) OK + SK + controls
Demonstrated knowledge (z)	−0.022 (0.049)		−0.058 (0.047)	−0.022 (0.037)
Subjective knowledge (z)		0.102** (0.041)	0.120*** (0.042)	0.091** (0.040)
Demographic controls	No	No	No	Yes
Parish fixed effects	No	No	No	Yes
Observations	826	803	803	802
— Lower panel: direct tests of the overconfidence configuration —				
SK on OK × opposition (interaction)	−0.020	(p = 0.52)	predicted	negative
Demonstrated knowledge on opposition	−0.002	(p = 0.96)	predicted	negative
Subjective knowledge on opposition	−0.085	(p = 0.045)	predicted	positive
Overconfidence on opposition	−0.079	(p = 0.055)	predicted	positive


Notes: Sample restricted to respondents expressing an opinion on the willingness items. Cluster-robust standard errors on 30 enumeration districts in parentheses. Columns (1)–(3) include no controls or fixed effects; column (4) adds both. Opposition is the negative of the standardized acceptance index. Lower-panel specifications include demographic controls and parish fixed effects. ** $p < 0.05$, *** $p < 0.01$. The final two cells of each lower-panel row state the sign the knowledge-overconfidence account predicts; the estimated signs are the opposite.

The lower panel provides the direct configuration tests. The interaction between demonstrated knowledge and opposition intensity in a regression for subjective knowledge—the specification closest to Fernbach et al. (2019)—is near zero (−0.020, $p = 0.52$). Regressing each measure on opposition, demonstrated knowledge is flat (−0.002, $p = 0.96$), subjective knowledge declines (−0.085, $p = 0.045$), and the overconfidence difference score also declines (−0.079, $p = 0.055$). The predicted signs for the latter two coefficients are positive. The point estimates therefore do not reproduce the published configuration; however, we do not characterize them as a definitive reversal because the difference-score result is only marginal and because our knowledge measure is a demonstrated-content index.

Figure 2 displays the same structure non-parametrically. Among complete opinion-holders, demonstrated knowledge shows no monotone gradient across the four opinion groups; subjective knowledge is also broadly flat except for a rise among strong supporters. The more pronounced separation is between complete respondents and those who provide at least one explicit ‘don’t know’ response: abstainers have mean demonstrated knowledge 0.38 standard deviations below the sample mean and subjective knowledge 0.36 standard deviations below it. Descriptively, the

dominant knowledge gradient therefore lies between substantive response and explicit uncertainty, not between opposition and support.

Figure 2. Objective and Subjective Knowledge across Response Groups



Points are group means of standardized knowledge with 95 per cent confidence intervals. Among complete opinion-holders neither measure exhibits the monotone gradient predicted by the knowledge-overconfidence account; respondents with at least one explicit 'don't know' response lie below every opinion-holding group on both measures.

Table 4 turns to the response margin. A one-standard-deviation increase in demonstrated knowledge is associated with an 8.7-percentage-point lower probability of abstaining in the bivariate specification and a 5.8-point lower probability once subjective knowledge, demographics, and parish fixed effects are included (analytic $p = 0.004$), against a base rate of 22.6 per cent. The wild-cluster bootstrap p -value is 0.008. Subjective knowledge has an independent association of similar direction (-0.051 , bootstrap $p = 0.003$). The result is therefore not confined to a mismatch between what respondents know and what they think they know: both lower demonstrated knowledge and lower self-assessed knowledge are associated with a greater propensity to answer 'don't know.'

The magnitude is large relative to the base rate. Across demonstrated-knowledge quintiles, abstention falls from roughly one third at the bottom to under one tenth at the top. This gradient remains descriptive: it combines any role of topic knowledge with correlated characteristics that are not fully observed. Logit average marginal effects (-0.061 for demonstrated knowledge and

−0.064 for subjective knowledge) are close to the linear estimates, and weighting by the household survey weight produces a demonstrated-knowledge coefficient of −0.080.

Table 4. Knowledge and Abstention

Dependent variable: abstains (0/1)	(1) OK only	(2) OK + SK	(3) + demographics	(4) + parish FE	(5) Logit AME
Demonstrated knowledge (z)	−0.087** *	−0.066***	−0.057***	−0.058***	−0.061***
	(0.017) {0.001}	(0.019)	(0.018)	(0.018) {0.008}	(0.014)
Subjective knowledge (z)		−0.058*** (0.015) {0.003}	−0.053*** (0.015)	−0.051*** (0.015) {0.003}	−0.064*** (0.016)
Demographic controls	No	No	Yes	Yes	Yes
Parish fixed effects	No	No	No	Yes	Yes
Observations	1,069	1,029	1,028	1,028	1,028

Notes: Linear probability models except column (5), which reports logit average marginal effects. Column (1) uses demonstrated knowledge only; columns (2)–(4) add subjective knowledge, demographic controls, and parish fixed effects sequentially; column (5) reports the logit average marginal effects for the full specification. Cluster-robust standard errors on 30 enumeration districts in parentheses; wild-cluster bootstrap p-values in braces (Rademacher weights, 1,999 replications, fixed seed; values are stable across seeds, ranging from 0.005 to 0.008 for objective knowledge and 0.002 to 0.003 for subjective knowledge). Base abstention rate 22.6 per cent. *** $p < 0.01$.

A natural concern is that the willingness result reflects something peculiar to that battery—for example, wording that makes ‘don’t know’ unusually attractive. The instrument permits a useful within-survey check because three additional batteries ask about perceived challenges of reuse in agriculture, tourism, and business, and each permits an explicit ‘don’t know’ response. The demonstrated-knowledge association appears in all three additional batteries, with coefficients larger in absolute value than for willingness: −0.143 for agriculture, −0.132 for tourism, and −0.166 for business, each with a bootstrap p-value of 0.001 (Table 5). It also predicts the total number of domains on which a respondent abstains, at −0.50 domains per standard deviation. These estimates are consistent in sign and magnitude across batteries, although the three sectoral questions share format and placement and should not be treated as statistically independent replications.

Two features of Table 5 matter for interpretation. First, subjective knowledge is not a consistent predictor across domains: it is negatively associated with abstention on personal willingness but is null or weakly positive in the three third-party batteries. Demonstrated topic knowledge is the more stable correlate. Second, willingness abstainers are largely different people from sectoral

abstainers: pairwise correlations are only 0.07, 0.08, and 0.09. This weighs against a simple story in which one fixed group of respondents declines every demanding question. The three sectoral batteries, however, are strongly correlated (0.52 to 0.75) and share format and placement.


Table 5. Abstention across Four Batteries

Battery	Abstain rate	Demonstrated knowledge	Subjective knowledge	Corr. with willingness
Willingness to use (Q611–614)	22.6%	−0.058 {0.007}	−0.051 (0.002)	—
Agriculture challenges (Q626)	37.6%	−0.143 {0.001}	+0.021 (0.297)	0.069
Tourism challenges (Q627)	58.4%	−0.132 {0.001}	+0.044 (0.028)	0.080
Business challenges (Q628)	57.1%	−0.166 {0.001}	+0.036 (0.055)	0.094
Number of domains abstained (0–4)	1.76 mean	−0.498 (0.000)	—	—

Notes: Each row is a separate regression with demographic controls and parish fixed effects, clustered on 30 enumeration districts. Wild-cluster bootstrap p-values in braces; analytic p-values in parentheses. The final column reports the raw correlation between each abstention indicator and abstention on the willingness battery. Correlations among the three sectoral batteries range from 0.52 to 0.75.

The appropriate claim is therefore consistency across two substantively distinct response contexts, personal willingness and third-party sectoral assessment, not four independent replications. Please, see Figure 3.

Figure 3. Abstention by Demonstrated-knowledge Quintile in Four Batteries




In every domain the share answering “don’t know” falls monotonically with demonstrated knowledge, though the batteries differ markedly in their base rates.

The index combines two item types, and they behave differently. Entered jointly, the recognition/content-production subscale predicts abstention strongly (-0.054 , $p = 0.005$), whereas the institutional-awareness subscale does not (-0.010 , $p = 0.53$). The association is therefore tied more closely to the ability to produce substantive content about reuse and water management than to awareness of a particular law, policy or infrastructure project. This pattern is useful because it reduces concern that the main result is driven entirely by local project awareness. It is not, however, a mechanism test: open-ended recognition items may also capture verbal fluency, numeracy or general response capacity.

The association is strongest in groups with higher baseline abstention. Among respondents with tertiary education the demonstrated-knowledge coefficient is -0.003 ($p = 0.89$; base abstention 13.4 per cent), compared with -0.078 ($p = 0.001$; base 27.2 per cent) among those without tertiary education. It is likewise larger among respondents aged sixty and over (-0.099 , $p < 0.001$; base 26.3 per cent) than among those under sixty (-0.026 , $p = 0.25$; base 19.7 per cent). Formal interactions are $+0.066$ ($p = 0.008$) for tertiary education and -0.059 ($p = 0.017$) for age sixty and over. These interactions show heterogeneity in the cross-sectional association in this sample. They do not establish that an information intervention would have larger treatment effects for less educated or older respondents. Because several subgroup contrasts are examined without multiplicity adjustment and the subgroup specification is not pre-registered, the heterogeneity exercise is best read as secondary, hypothesis-generating evidence about where explicit uncertainty is concentrated. Please, see Figure 4.

Association between Demonstrated Knowledge and Abstention, by Subgroup



Points are subgroup-specific coefficients in percentage points with 95 per cent cluster-robust confidence intervals, each from a separate regression controlling for subjective knowledge and age. These are secondary heterogeneity analyses.

7. Robustness


This section examines threats to the two descriptive claims: demonstrated topic knowledge is close to unrelated to the direction of stated willingness among complete opinion-holders but is strongly associated with explicit ‘don’t know’ response. We focus on measurement, response capacity, selection into the opinion-holder sample, and the precision of the conditional direction estimate. Figure 5 summarizes the demonstrated-knowledge association across specifications.³ The main coefficient (-0.058) is close to the bivariate value (-0.087), the logit average marginal effect (-0.061), the survey-weighted estimate (-0.080), and an alternative standardized mean-of-z-scores index (-0.054). The relationship also appears item by item: demonstrated knowledge predicts ‘don’t know’ on each of the four willingness items separately, with coefficients from -0.037 to -0.057 and all p-values below 0.02. Dropping each knowledge item in turn leaves the coefficient between -0.040 and -0.067 (maximum $p = 0.039$). A stricter outcome -‘don’t know’ on all four willingness items- remains negatively associated with demonstrated knowledge (-0.034 , $p < 0.001$) against a 7.3 per cent base rate.

³ Table A1 in the Appendix reports the full set.

A further check addresses coding of the demonstrated-knowledge index. Scoring multiple-response items requires judgments about which answers count as substantive, and the baseline rules do not treat every ‘other’ category identically. We therefore re-estimate under four alternative coding rules: counting ‘other’ everywhere, excluding it everywhere, adding the broad environmental option, and a fully harmonized rule. The resulting indices correlate between 0.98 and 0.99 with the baseline and have similar reliability (alpha 0.695 to 0.711). The response-margin coefficient ranges from -0.054 to -0.064 and is statistically significant under every rule ($p \leq 0.007$), while the conditional direction coefficient ranges from -0.013 to -0.038 and is statistically insignificant under every rule ($p \geq 0.30$).

The substantive contrast between the two margins is therefore not driven by a particular coding convention. The direction regressions use only respondents with no ‘don’t know’ response on any willingness item, so the acceptance index is complete for every observation used.⁴ This separation removes one concern, averaging an incomplete scale, but not the broader selection issue addressed below.

Figure 5. Abstention Results across Specifications



Each point is the estimated change in the probability of abstaining associated with a one-standard-deviation increase in demonstrated knowledge, with 95 per cent cluster-robust confidence intervals. The main specification corresponds to column (4) of Table 4; the content-free nonresponse specification is discussed below. Intervals use the t critical value with 29 degrees of freedom.

The most serious interpretive threat is that the demonstrated-knowledge index partly proxies a general capacity or willingness to produce survey answers—literacy, numeracy, verbal fluency, or

⁴ The 990-household acceptance index reported in Table 2 includes partial abstainers and is used only for descriptive statistics.

cognitive engagement. This concern is especially salient because the strongest subscale consists of recognition/content-production items. If the same latent response capacity drives both the knowledge index and ‘don’t know’ outcomes, a purely reuse-specific interpretation would be overstated. The data support part of this concern. Respondents with higher reuse knowledge are also less likely to leave an out-of-module household-income item unresolved, and demonstrated knowledge predicts an out-of-module nonresponse index (-0.076 , $p < 0.001$). Decomposing the income item is informative: the strongest association is with being unable to estimate income (-0.110 , $p < 0.001$) and not privacy refusal (-0.027 , $p = 0.060$) or a generic ‘don’t know’ answer (-0.012 , $p = 0.21$).⁵

We then include the out-of-module nonresponse index directly in the main abstention regression. The demonstrated-knowledge coefficient attenuates from -0.058 to -0.047 and remains statistically significant ($p = 0.015$), a reduction of about 19 per cent in absolute magnitude. This shows that measured general response capacity explains some of the covariance while leaving a residual association. It is not valid to interpret the 19 per cent attenuation as a causal decomposition: the control is an imperfect proxy and could itself be correlated with unobserved determinants of both variables. The defensible conclusion is narrower—the main result is not wholly explained by measured out-of-module nonresponse, but neither is it a pure fact about reuse knowledge.

The response-margin result creates a selection problem for the direction regression. If higher-knowledge respondents are more likely to provide a complete willingness profile, then the observed opinion-holder sample is not a random slice of the population. A zero conditional association could therefore differ from the population relationship that would be observed if every respondent had an observed or imputable acceptance value.⁶

The fully specified demonstrated-knowledge coefficient among complete opinion-holders is -0.022 with a 95 per cent confidence interval of $[-0.094, 0.049]$ standard deviations. The interval

⁵ This pattern is consistent with a shared numeracy or response-capacity component, although the income question is not a pure response-capacity measure and may also reflect sensitivity or privacy.

⁶ We illustrate sensitivity with two common-value imputations. Table A3 in the Appendix reports the exercise. Assigning every abstainer the maximum acceptance value yields a demonstrated-knowledge coefficient of -0.084 ; assigning every abstainer the minimum yields $+0.118$. These assignments are deliberately polar and easy to interpret, but they are not sharp or worst-case bounds on the regression coefficient. The exercise therefore shows that the conditional estimate is not invariant to assumptions about abstainers, while leaving the population-wide latent association unidentified.

rules out a large negative linear association and also excludes positive associations above about five hundredths of a standard deviation in this selected sample. This is stronger than saying that the coefficient is statistically insignificant: the estimate is centered close to zero with reasonably tight uncertainty. It does not, however, reject nonlinear, subgroup-specific, or causal mechanisms linking knowledge to attitudes, nor does it resolve selection into the opinion-holder sample.

Five limitations bound the interpretation. First, the design is a single cross-section. Knowledge may shape opinion formation, opinion may motivate information acquisition, and both may reflect unobserved traits; all reported coefficients are therefore associations. Second, an explicit ‘don’t know’ response is an imperfect and behaviorally mixed indicator. It can represent recognized uncertainty, satisficing, reluctance to guess, or item-specific difficulty. The low knowledge of abstainers and the replication across substantively distinct batteries make the response margin informative, but they do not identify its latent meaning. Third, the demonstrated-knowledge index is not a pure factual test. Several items require respondents to produce substantive content, so it may partly capture verbal or cognitive response capacity. provides evidence that this concern is empirically relevant. Fourth, the direction analysis conditions on complete response, leaving population-wide acceptance–knowledge relationships unidentified for abstainers. Fifth, the parent evaluation’s two community arms were formed purposively. We pool them, control for arm membership and parish, and draw no inference from the arm contrast; the demonstrated-knowledge coefficient is similar in the two arms (-0.061 in each), which is reassuring about geographic concentration but does not make the sample nationally representative. External validity is necessarily limited: Barbados is a small, water-scarce island at a particular stage of reuse policy development, and the response margin may shrink once an issue becomes familiar, politicized, or routinely experienced. The paper’s strongest general claim is therefore a boundary condition and a measurement lesson, not a universal model of water-reuse opinion.

8. Why Might Water Reuse Differ from Genetically Modified Foods?

Why might the knowledge-overconfidence configuration organize opposition to, for instance, genetically modified foods but not appear here? The cross-section cannot adjudicate mechanisms, so the following explanations are hypotheses disciplined by the observed pattern. They matter because each suggests a different reason why knowledge might separate substantive responders from abstainers without strongly sorting the direction of opinion.

One possibility concerns the nature of the aversion. Some opposition to genetically modified foods may be articulated through claims about hidden biological harms, an epistemic domain in which respondents can perceive themselves as informed. Water-reuse aversion often has a strong affective component rooted in contamination and disgust (Rozin et al., 2015; Russell and Lux, 2009). A respondent can experience that reaction without claiming technical knowledge of treatment processes. If overconfidence is most likely when opposition is expressed through claims of understanding, affect-heavy domains may generate a weaker objective–subjective knowledge signature even when opposition is real.

A second possibility is politicization and identity loading. Light et al. (2022) show that the overconfidence pattern varies across issues, and related work on scientific disagreement emphasizes social identity and epistemic deference. Our survey does not directly measure whether water reuse in Barbados is tied to partisan, religious, or other group identities, so we cannot assert that it is unpoliticized. The empirical pattern is nevertheless consistent with a domain in which many respondents have not received strong identity cues that would make defending a position socially valuable. A third possibility concerns the stage of issue development. Genetically modified foods have been publicly contested for decades in many settings. Reuse in Barbados appears less familiar to many households, as suggested by the low demonstrated-knowledge scores and large explicit ‘don’t know’ shares. In an early stage of issue development, information may primarily separate respondents who can formulate a judgment from those who cannot; only later, as arguments, identities, and advocacy networks develop, might knowledge and confidence sort the direction and intensity of positions. The present data are consistent with that sequence but cannot establish it without repeated measurement.

These explanations are not mutually exclusive. Distinguishing them requires longitudinal data as a reuse program becomes more salient, experimental variation in information and source cues, or comparable surveys across domains that differ systematically in affective intensity, identity loading, and issue maturity. Such designs would also reveal whether the response-margin association documented here is a temporary feature of early issue development or a persistent characteristic of reuse attitudes.

9. Conclusions

Three implications follow, but they should be read as diagnostic implications of an observational study and not as estimates of policy effects. First, the key descriptive fact is the size and composition of the explicit-uncertainty margin. Twenty-two point six per cent of respondents give at least one ‘don’t know’ answer on personal willingness, and between 37.6 and 58.4 per cent do so on the sectoral batteries. These respondents have, on average, lower demonstrated and self-assessed knowledge. This identifies a plausible audience for communication and deliberation, but it does not establish that less educated or older respondents will respond more strongly to an intervention; the subgroup coefficients are associations, not treatment effects. Second, the two-part framework yields a testable prediction for future evaluation. If information primarily operates on the response margin, an effective intervention should first reduce explicit ‘don’t know’ responses and increase the share of respondents able or willing to state a directional judgment. The direction of newly articulated opinions is not predetermined. A randomized information experiment or panel survey should therefore report both changes in abstention and changes in acceptance among substantive responders, instead of treating a decline in ‘don’t know’ as synonymous with movement toward support. Third, the results argue against deleting explicit ‘don’t know’ responses before modeling public acceptance. Where that share is large, a two-part analysis of response existence and response direction can be more informative than a single attitude mean computed only on substantive answers.

The central empirical result is not that knowledge determines support or opposition. Among complete opinion-holders, demonstrated knowledge is close to unrelated to acceptance, while perceived knowledge is positively associated with acceptance. The strongest and most robust knowledge gradient lies on the response margin: respondents with lower demonstrated knowledge are substantially more likely to answer ‘don’t know,’ and the same relationship appears in sectoral assessments. The association survives alternative estimators and coding rules and remains after controlling for an out-of-module response measure, although that control attenuates the coefficient. These patterns are consistent with recognized uncertainty playing a meaningful role, while the cross-sectional design and the content-production nature of the knowledge index prevent a purely causal or purely knowledge-specific interpretation.

The broader contribution is to treat public opinion on an unfamiliar environmental technology as having two observable margins: whether a respondent supplies a substantive response and, conditional on doing so, the direction of that response. In these data, demonstrated knowledge is much more strongly associated with the first. The finding qualifies the generality of the confidently mistaken-opponent account, connects water-reuse research to the survey literature on response uncertainty, and shows why explicit ‘don’t know’ responses contain information that is lost when they are deleted.

References

- Al-Saidi, M. (2021). From acceptance snapshots to the social acceptability process: Structuring knowledge on attitudes towards water reuse. *Frontiers in Environmental Science* 9, 633841.
- Cameron, A. C., and D. L. Miller (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Converse, J. M. (1976). Predicting no opinion in the polls. *Public Opinion Quarterly* 40(4), 515–530.
- Dolnicar, S., A. Hurlimann, and B. Grün (2011). What affects public acceptance of recycled and desalinated water? *Water Research* 45(2), 933–943.
- Elkjær, M. A., and C. Wlezien (2025). Estimating public opinion from surveys: The impact of including a 'don't know' response option in policy preference questions. *Political Science Research and Methods* 13(3), 663-679.
- Fernbach, P. M., N. Light, S. E. Scott, Y. Inbar, and P. Rozin (2019). Extreme opponents of genetically modified foods know the least but think they know the most. *Nature Human Behaviour* 3(3), 251–256.
- Fielding, K. S., and A. H. Roiko (2014). Providing information promotes greater public support for potable recycled water. *Water Research* 61, 86-96.
- Haener, M. K., and W. L. Adamowicz (1998). Analysis of 'don't know' responses to referendum contingent valuation questions. *Agricultural and Resource Economics Review* 27(2), 218-230.
- Hurlimann, A., and S. Dolnicar (2016). Public acceptance and perceptions of alternative water sources: A comparative study in nine locations. *International Journal of Water Resources Development* 32(4), 650-673.
- Krosnick, J. A. (1991). Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology* 5(3), 213–236.
- Krosnick, J. A., A. L. Holbrook, M. K. Berent, R. T. Carson, W. M. Hanemann, R. J. Kopp, R. C. Mitchell, S. Presser, P. A. Ruud, V. K. Smith, W. R. Moody, M. C. Green, and M. Conaway (2002). The impact of 'no opinion' response options on data quality: Non-attitude reduction or an invitation to satisfice? *Public Opinion Quarterly* 66(3), 371-403.
- Kruger, J., and D. Dunning (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology* 77(6), 1121–1134.
- Kuha, J., S. Butt, M. Katsikatsou, and C. J. Skinner (2018). The effect of probing 'don't know' responses on measurement quality and nonresponse in surveys. *Journal of the American Statistical Association* 113(521), 26-40.
- Light, N., P. M. Fernbach, N. Rabb, M. V. Geana, and S. A. Sloman (2022). Knowledge overconfidence is associated with anti-consensus views on controversial scientific issues. *Science Advances* 8(29), eabo0038.

- Luskin, R. C., and J. G. Bullock (2011). “Don’t know” means “don’t know”: DK responses and the public’s level of political knowledge. *Journal of Politics* 73(2), 547–557.
- Menegaki, A. N., N. Hanley, and K. P. Tsagarakis (2007). The social acceptability and valuation of recycled water in Crete: A study of consumers’ and farmers’ attitudes. *Ecological Economics* 62(1), 7–18.
- Moesker, K., U. Pesch, and N. Doorn (2024). Public acceptance in direct potable water reuse: A call for incorporating responsible research and innovation. *Journal of Responsible Innovation* 11(1), 2304382.
- Nkhoma, P. R., K. Alsharif, E. Ananga, M. Eduful, and M. Acheampong (2021). Recycled water reuse: What factors affect public acceptance? *Environmental Conservation* 48, 278-286.
- Ormerod, K. J., and C. A. Scott (2013). Drinking wastewater: Public trust in potable reuse. *Science, Technology, & Human Values* 38(3), 351-373.
- Po, M., J. D. Kaercher, and B. E. Nancarrow (2003). Literature review of factors influencing public perceptions of water reuse. CSIRO Land and Water Technical Report 54/03.
- Roodman, D., M. Ø. Nielsen, J. G. MacKinnon, and M. D. Webb (2019). Fast and wild: Bootstrap inference in Stata using boottest. *Stata Journal* 19(1), 4–60.
- Roopnarine, R., K. Baird, M. Hosein, R. Jackson, S. Salim, A. Cephas, O. Daley, S. Gangapersad, S.-J. Govia, A. Cashman, A. Crichlow, and F. Pinongcos (2023). Integrating wastewater reuse into water management schemes of Caribbean SIDS: A Trinidad and Tobago case study. *Water Policy* 25(12), 1161-1174.
- Ross, V. L., K. S. Fielding, and W. R. Louis (2014). Social trust, risk perceptions and public acceptance of recycled water: Testing a social–psychological model. *Journal of Environmental Management* 137, 61–68.
- Rozin, P., B. Haddad, C. Nemeroff, and P. Slovic (2015). Psychological aspects of the rejection of recycled water: Contamination, purification and disgust. *Judgment and Decision Making* 10(1), 50–63.
- Russell, S., and C. Lux (2009). Getting over yuck: Moving from psychological to cultural and sociotechnical analyses of responses to water recycling. *Water Policy* 11(1), 21–35.
- Sardana, P., A. Javernick-Will, and S. M. Cook (2025). Facilitators and barriers of global water reuse: A systematic literature review. *ACS ES&T Water* 5(1), 3–19.
- Tsagarakis, K. P., and N. Georgantzis (2003). The role of information on farmers’ willingness to use recycled water for irrigation. *Water Science and Technology: Water Supply* 3(4), 105–113.
- Wang, H. (1997). Treatment of 'don't-know' responses in contingent valuation surveys: A random valuation model. *Journal of Environmental Economics and Management* 32(2), 219-232.
- Wang, Y., S. Mehrabi, and K. Ikuma (2026). Acceptance of water reuse adoption in small communities: Analysis of driving factors and policy preferences. *Environmental Research* 300, 124439.
- World Bank (2023). Renewable internal freshwater resources per capita (cubic meters). *World Development Indicators*. Washington, DC: World Bank.

Appendix. Supplementary Results

Table A1. Robustness of the Abstention Association

Specification	Coefficient	Std. error / p	Observations
Main (SK, demographics, parish FE)	−0.058	(0.018)	1,028
Bivariate	−0.087	(0.017)	1,069
Logit, average marginal effect	−0.061	(0.014)	1,028
Household survey weights	−0.080	p < 0.001	1,028
Mean-of-z knowledge index	−0.054	p = 0.006	1,028
Recognition subscale only	−0.056	p = 0.004	1,028
Strict abstention (all four items)	−0.034	p < 0.001	1,028
Control for out-of-module nonresponse	−0.047	p = 0.015	1,028
Leave-one-out range (10 indices)	−0.040 to −0.067	max p = 0.039	1,028
Alternative coding rules (4 variants)	−0.054 to −0.064	max p = 0.007	1,028
Item-by-item (Q611–Q614)	−0.037 to −0.057	all p < 0.02	1,028
Treatment arm only	−0.061	p = 0.057	493
Comparison arm only	−0.061	p = 0.022	535

Notes: Each row is the coefficient on standardized demonstrated knowledge in a linear probability model for abstention, unless otherwise stated. Cluster-robust standard errors on 30 enumeration districts. Wild-cluster bootstrap p = 0.008 for the main specification.

Table A2. Heterogeneity in the Knowledge–Abstention Association

Subgroup	Coefficient	p-value	Base abstention
No tertiary education	−0.078	0.001	27.2%
Tertiary education	−0.003	0.894	13.4%
Age 60 and over	−0.099	<0.001	26.2%
Under 60	−0.026	0.251	19.7%
Female respondent	−0.076	0.001	22.0%
Male respondent	−0.050	0.065	23.4%
Farming household	−0.053	0.064	20.7%
Non-farming household	−0.064	0.004	23.0%

Notes: Separate regressions by subgroup, controlling for subjective knowledge and age, clustered on 30 enumeration districts. Formal interaction terms: objective knowledge x tertiary education = +0.066 (p = 0.008); objective knowledge x age 60+ = −0.059 (p = 0.017). These are secondary heterogeneity analyses of associations, not treatment-effect estimates; no multiplicity adjustment is applied.

Table A3. Selection Sensitivity Under Common-Value Imputations

Assignment of abstainers	Coefficient	p-value	N
Excluded (observed, opinion-holders only)	−0.022	0.549	802
All assigned maximum acceptance	−0.084	0.015	1,027
All assigned minimum acceptance	+0.118	0.050	1,027