CSEF

Centre for Studies in Economics and Finance

WORKING PAPER NO. 793

Observability and Social Interactions in Consumption

Tullio Jappelli and Luigi Pistaferri

July 2026

University of Naples Federico II

University of Salerno

Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS AND STATISTICS – UNIVERSITY OF NAPLES FEDERICO II
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unina.it
ISSN: 2240-9696

CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 793

Observability and Social Interactions in Consumption

Tullio Jappelli* **and Luigi Pistaferri†**

Abstract
This paper examines how social context affects consumption responses to transitory income shocks. Using a randomized survey experiment embedded in the Italian Survey of Consumer Expectations, we vary whether an unexpected €1,000 employer-specific bonus is observable to coworkers and whether it is received only by the respondent or by all coworkers. The design separately identifies the effect of observability for an idiosyncratic bonus and the incremental effect of making a publicly observable bonus common within the workplace, a well-defined social context. Making an idiosyncratic bonus publicly observable reduces the MPC by about 3 percentage points relative to a private idiosyncratic bonus. Conditional on observability, making the bonus common among coworkers raises the MPC by about 6 percentage points. Interpreted through a linear social-interactions model, the latter effect implies a social-interaction parameter of approximately 0.16 and a social multiplier of about 1.2. These findings suggest that aggregate consumption responses to fiscal policy depend not only on individual behavior but also on the social context in which income changes occur.

Keywords: Marginal Propensity to Consume; Social Interactions; Income Shocks.

JEL Classification: D12, D14, D15, C8, C99.

Acknowledgments: We thank Orazio Attanasio, Michalis Haliassos and seminar participants at the ECB Household Finance and Consumption research seminar. The research was funded by the European Union - Next Generation EU, part of the GRINS - Growing Resilient, Inclusive and Sustainable Project framework (GRINS PE00000018 – CUP E63C22002140007). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union.

* University of Naples Federico II, CSEF and CEPR. E-mail: tullio.jappelli@unina.it

† Stanford University, SIEPR, NBER and CEPR.

1. Introduction

The marginal propensity to consume (MPC) plays a central role in macroeconomics because it can be used to evaluate the impact of fiscal policy. In standard consumption models, households smooth consumption intertemporally in response to expected lifetime resources, with consumption choices determined primarily by preferences, lifetime resources, and expectations (including liquidity constraints). Modern heterogeneous-agent models emphasize that there might be substantial variation in MPCs due to households differing in resources, asset composition, borrowing constraints, and income risk.

At the same time, a long tradition in economics emphasizes that individuals derive utility not only from their own consumption, but also from social status and from comparing their consumption to that of their peers (Veblen, 1899; Duesenberry, 1949). Recent papers have shown that households adjust expenditures in response to neighbors' or peers' consumption, particularly for so-called "visible goods" (Charles et al., 2009; Kuhn et al., 2011; De Giorgi et al. 2020). These mechanisms imply that fiscal shocks may propagate through social networks, amplifying or dampening aggregate responses beyond what standard models predict.

This paper studies two distinct channels through which social context may affect consumption responses to income shocks: observability and peer interactions. Observability captures whether members of the reference group can observe an individual's idiosyncratic income change, which may activate reputational concerns, social norms, or pressures towards restraint or conformity. Peer effects capture the possibility that, as the consumption of peers change, own consumption may also change because of "keeping up with the Joneses" effects. Distinguishing between these channels is important because they may have opposite aggregate consequences. Visible, idiosyncratic income shocks may induce restraint if individuals fear social scrutiny or norm violations, whereas peer effects may increase aggregate spending through a social multiplier effect channel.

We study the importance of social mechanisms using a randomized survey experiment in which individuals allocate a hypothetical €1,000 unexpected employer-specific bonus to consumption, saving, and transfers to family or friends. The research design varies whether the bonus is publicly observable and whether it is received only by the respondent or by all coworkers. It therefore separately identifies the effect of observability for an idiosyncratic bonus and the incremental effect of making a publicly observable bonus common within the workplace. The workplace constitutes a particularly relevant environment because coworkers

are salient peers who frequently share information, compare outcomes, and observe each other's spending, especially on visible goods such as clothing, personal services, jewelry, and transportation. Our findings reveal that the social context affects consumption responses. Observability reduces the marginal propensity to consume by roughly 3 percentage points, consistent with the existence of norm-based restraint or reputational concerns discouraging visible spending following a private windfall. In contrast, when all coworkers receive the bonus, the MPC increases by approximately 6 percentage points, suggesting the presence of strategic complementarities and peer amplification. Interpreted through our linear social-interactions model and under its maintained assumptions, the estimate implies a social-interaction parameter of approximately 0.16 and a social multiplier of about 1.2. Under this interpretation, common shocks generate stronger spending responses because individuals respond to the higher anticipated consumption of their peers. These findings imply that social interactions can either attenuate or amplify fiscal policy interventions depending on the structure of the shock and the surrounding social environment.

More broadly, the paper suggests that models of household consumption and fiscal policy may need to incorporate the social context more explicitly. If MPCs depend not only on household characteristics but also on whether income shocks are publicly observable or shared within a reference group, then the aggregate effects of fiscal interventions may vary systematically across institutional and social environments. Common and salient transfers may generate larger aggregate demand effects than isolated or private transfers, because they may activate peer amplification. Conversely, public observability may dampen spending responses when social norms favor prudence or restraint. Understanding these mechanisms is therefore important not only for household finance and behavioral economics, but also for the design and evaluation of stabilization policy. Our results also suggest that MPCs elicited from purely private, idiosyncratic shocks can misstate the consumption response to fiscal interventions whose social structure differs from that benchmark. In our design, they understate the response to broad transfers received by the whole reference group (a net effect of 2.5 percentage points). By contrast, because the observability effect is negative, they may overstate the response to transfer that are is publicly observable by only a minority of the group.

The paper contributes to three broad research areas. First, to research on social multipliers and aggregate amplification. Theoretical work has long emphasized that interactions among agents can magnify the effects of individual shocks through complementarities in behavior

(Becker and Murphy, 2000; Glaeser et al., 2003). In macroeconomics, social multipliers have been studied in contexts such as labor supply, education, technology adoption, and local demand spillovers. More recently, De Giorgi et al. (2020) argue that consumption responses may also exhibit network amplification when households observe and react to the behavior of others (in their context, coworkers). Yet direct experimental evidence on the magnitude of spillovers in consumption remains scarce. By randomly varying whether income shocks are idiosyncratic or common within a workplace peer group, we provide an experimental estimate of the effect of shock commonality on consumption responses. Interpreted through our model, this reduced-form effect maps into a peer-interaction parameter and an implied social multiplier.

Second, the paper contributes to the literature on peer effects and social interactions. Identifying peer effects is notoriously difficult because of correlated unobservable variables and the presence of the "reflection problem" (Manski, 1993). A large empirical literature has attempted to overcome these challenges using randomized assignment, natural experiments, or directly modeling the network structure (Sacerdote, 2001; Carrell et al., 2009). A more recent strand uses field experiments to identify the role of observability and peer behavior in specific decision domains, including educational investments (Bursztyn and Jensen, 2015) and financial choices (Bursztyn et al., 2014, 2018). In the consumption literature, the available evidence suggests that households respond to the spending behavior of neighbors or coworkers, especially for goods and housing that are more "visible" or publicly observable (Bertrand and Morse, 2016; De Giorgi et al., 2020). Our setting provides a natural and economically meaningful reference group (coworkers), and to our knowledge this is the first paper to experimentally vary the social context of a transitory income shock (its observability and its commonality within the reference group), thus bridging the elicited-MPC literature (Jappelli and Pistaferri, 2014; Fuster et al., 2021) with the observational consumption-peer-effects literature (Bertrand and Morse, 2016; De Giorgi et al., 2020).

Finally, the paper contributes to two related methodological literatures: the measurement of MPCs through hypothetical survey questions (Jappelli and Pistaferri, 2020; Fuster et al., 2021) and the broader use of survey experiments in macroeconomics and public finance (Stantcheva, 2023). Hypothetical MPC questions are widely used to recover consumption responses to shocks that mirror fiscal policy interventions under controlled conditions, but they typically abstract from the social environment and implicitly assume that income shocks are privately observed. We show that introducing social context into these hypothetical scenarios

changes measured MPCs, suggesting that conventional elicitation methods may overlook an important source of heterogeneity in consumption responses (or may implicitly embed it, depending on how private the hypothetical shock is perceived to be). Our design extends the survey experiment methodology to social interactions in consumption by manipulating both observability and commonality of shocks within a realistic workplace environment.

The remainder of the paper is organized as follows. Section 2 uses a simple model of consumption with observability and social interactions to derive an estimating equation that matches the type of data collected in our randomized experiment. Section 3 describes the survey experiment and the empirical design. Section 4 reports the main results on the effects of observability and shared shocks on consumption responses. Section 5 explores heterogeneity in the estimated effects across individuals and workplace environments. Section 6 concludes and discusses the broader implications for consumption behavior and fiscal policy.

2. A model of social interactions

We adopt a parsimonious model of consumption with social interactions that nests both peer effects and observability, while allowing for individual income shocks. The framework derives a linear-in-means model of social behavior (Manski, 1993; Blume et al., 2015) based on preferences that allow for "keeping up with the Joneses" effects (Abel, 1990) as well as a utility shifter from observability of income shocks.

Consider an individual i who receives a transitory income shock Δy_i (a bonus) from her employer. She allocates the bonus between current consumption and saving, taking into account three forces: the intertemporal tradeoff (consuming today reduces future resources), social comparison with the consumption of a reference group of coworkers, and the reputational or normative consequences of spending when the income gain is publicly observable. We model this as a standard intertemporal problem in which current utility depends on own consumption, the average consumption of peers, and the publicly observable component of the shock.[1]

[1] The model abstracts from transfers to family and friends, although the empirical question allows respondents to allocate a hypothetical bonus among consumption, saving and transfers. This can be interpreted as a two-stage budgeting problem. In the first stage, households earmark part of the bonus for transfers, based on altruistic, family or reciprocity considerations. In the second stage, they allocate the remaining resources between own consumption and saving, taking into account the workplace social environment. Under this interpretation, if transfer decisions are unaffected by the workplace treatments, the model applies to the second-stage consumption-saving margin. Consistent with this assumption, we find no detectable effect of either observability or the common-shock treatment on the marginal propensity to transfer. The evidence therefore suggests that the treatment induce reallocations primarily along the consumption–saving margin rather than through transfers to family and friends.

Let $\bar{c}_{g,-i}$ denote the average consumption level of i's relevant peer group, excluding individual i, and superscript 0 the no-bonus baseline. In the bonus scenario, income changes from y_i^0 to y_i, so that the income shock is $\Delta y_i = y_i - y_i^0$. The corresponding consumption deviations from baseline are: $\Delta c_i = c_i - c_i^0$, and $\Delta \bar{c}_{g,-i} = \bar{c}_{g,-i} - \bar{c}_{g,-i}^0$. Finally, let P_i an indicator equal to one when the income shock is publicly observable, and call $z_i = P_i \Delta y_i$ the publicly observable component of the income shock.

Individual i chooses current consumption c_i to maximize

$$U = u(c_i, \bar{c}_{g,-i}) + S(c_i, z_i) + \beta V(A_{i1}) \tag{1}$$

subject to the constraint:

$$A_{i1} = (1 + r)A_{i0} + y_i - c_i.$$

Here $u(c_i, \bar{c}_{g,-i})$ is current utility, including any direct dependence on peers' consumption. The value function $V(A_{i1})$ is the continuation value as a function of next-period assets and captures the intertemporal budget tradeoff: consuming more today reduces future resources. The term $S(c_i, z_i)$ is a social or reputational payoff from spending when the income shock is observable. We normalize this term to $S(c_i, 0) = 0$, for all c_i. Thus, when the shock is not publicly observable, the social payoff term disappears. Note that this implies $S_c^0 = S_{cc}^0 = 0$.

The no-bonus baseline satisfies: $A_{i1}^0 = (1 + r)A_{i0} + y_i^0 - c_i^0$. Subtracting the baseline budget constraint from the bonus-scenario budget constraint yields: $A_{i1} = A_{i1}^0 + \Delta y_i - \Delta c_i$. The general first-order condition for consumption is:

$$u_c\left(c_i, \bar{c}_{g,-i}\right) + S_c(c_i, z_i) - \beta V_A(A_{i1}) = 0 \tag{2}$$

At the baseline interior optimum,

$$u_c^0 = \beta V_A^0,$$

where all derivatives with superscript 0 are evaluated at (c_i^0,$\bar{c}_{g,-i}^0$,$z_i = 0$, and A_{i1}^0) and we have used the fact that $S_c^0 = 0$.

Linearizing this first-order condition around the no-bonus baseline gives

$$0 \cong [u_c^0 + u_{cc}^0\,\Delta c_i + u_{c,\bar{c}}^0 \Delta \bar{c}_{g,-i}] + [S_c^0 + S_{cc}^0\,\Delta c_i + S_{cz}^0\,P_i \Delta y_i] - \beta[V_A^0 + V_{AA}^0(\Delta y_i - \Delta c_i)].$$

Using the fact that $S_c^0 = S_{cc}^0 = 0$ and that $u_c^0 = \beta V_A^0$, this becomes:

$$(u_{cc}^0 + \beta V_{AA}^0)\Delta c_i + u_{c,\bar{c}}^0\,\Delta \bar{c}_{g,-i} + (S_{cz}^0 P_i - \beta V_{AA}^0)\Delta y_i = 0. \tag{3}$$

Rearranging, it yields:

$$\Delta c_i = \delta \Delta y_i + \alpha \Delta \bar{c}_{g,-i} + \phi P_i \Delta y_i, \tag{4}$$

where $\delta = \frac{\beta V_{AA}^0}{u_{cc}^0 + \beta V_{AA}^0}$, $\alpha = -\frac{u_{c,\bar{c}}^0}{u_{cc}^0 + \beta V_{AA}^0}$, $\phi = -\frac{S_{cz}^0}{u_{cc}^0 + \beta V_{AA}^0}$.

The parameter δ is the marginal propensity to consume with respect to a private, idiosyncratic transitory shock, i.e., the standard MPC (defined as $MPC_i = \frac{\Delta c_i}{\Delta y_i}$). If preferences are concave, there is diminishing marginal utility of consumption and diminishing marginal utility of wealth, implying $u_{cc}^0 < 0$, $V_{AA}^0 < 0$, and hence $0 < \delta < 1$.

The parameter α captures social interactions in consumption. If peer consumption and own consumption are complements, peer consumption raises the marginal utility of own consumption ($u_{c,\bar{c}}^0 > 0$), implying $\alpha > 0$.[2] In this case individual consumption increases with peers' consumption, capturing mechanisms such as the "keeping up with the Joneses" or conspicuous consumption ("status-seeking" behavior).

The parameter ϕ captures the effect of shock observability on spending. The numerator of ϕ can be either positive or negative, as public observability of the shock may either increase or reduce the social value of consumption (S_{cz}^0). The overall expression for the consumption effect of public income is therefore ambiguous. When $\phi > 0$, observability raises the MPC, consistent with signaling or status-seeking motives. When $\phi < 0$, observability lowers the

[2] Throughout, we restrict attention to preferences satisfying: $u_{c,\bar{c}}^0 \geq 0$ (consumption complementarity in the sense of "keeping up with the Joneses") and $u_{c,\bar{c}}^0 < |u_{cc}^0| + \beta|V_{AA}^0|$. The first inequality ensures $\alpha \geq 0$; the second ensures $\alpha < 1$ and hence a finite social multiplier. Standard external-habit preferences $U(c - \gamma\bar{c})$ with $\gamma \in [0, 1)$ satisfy both conditions (Abel, 1990).

MPC, reflecting social norms, reputational concerns, or pressures toward moderation.[3] This makes the effect of visibility an empirical question.

Our survey experiment assigns respondents to three scenarios that identify these parameters separately.[4] In the first scenario, the bonus is private and idiosyncratic ($P_i = 0$ and the peers are unaffected: $\Delta\bar{c}_{g,-i} = 0$). In this case, the MPC is: $MPC_{P=0,All=0} = \frac{\Delta c_i}{\Delta y_i} = \delta$. In the second scenario, the bonus is publicly observable but still idiosyncratic ($P_i = 1$ but the peers are still unaffected ($\Delta\bar{c}_{g,-i} = 0$). Since $MPC_{P=1,All=0} = \frac{\Delta c_i}{\Delta y_i} = \delta + \phi$; the difference from the first scenario identifies the observability effect. In the third scenario, the bonus is publicly observable and received by all coworkers. $\Delta\bar{c}_{g,-i} = \Delta c_i, P_i = 1$.[5] In this case, $\Delta c_i = \delta\Delta y_i + \alpha\Delta c_i + \phi\Delta y_i$. Under a symmetric equilibrium, the average change in peers' consumption equals the own consumption change ($\Delta\bar{c}_{g,-i} = \Delta c_i$), and $MPC_{P=1,All=1} = \frac{\Delta c_i}{\Delta y_i} = \frac{\delta+\phi}{1-\alpha}$.

Hence, the general expression for the MPC that covers the three scenarios is:

$$MPC_i = MPC_{P=0,All=0}(1 - P_i)(1 - All_i) + MPC_{P=1,All=0}P_i(1 - All_i) + MPC_{P=1,All=1}P_iAll_i,$$

which can be rewritten (since in our experiments $P_iAll_i = All_i$) as:

$$MPC_i = \beta_1 + \beta_2 P_i + \beta_3 All_i + X_i'\gamma + \epsilon_i \quad (5)$$

where $\beta_1 = \delta, \beta_2 = \phi$, and $\beta_3 = \frac{\alpha(\delta+\phi)}{1-\alpha}$, All_i is an indicator for the common-shock scenario, and we have added observable socio-demographic characteristics and unobserved heterogeneity in preferences.

The reduced form parameter β_1 identifies the MPC when the change in income is idiosyncratic and private (δ); the reduced form parameter β_2 identifies the effect of

[3] An alternative mechanism is that, when the individual is the only recipient of a public bonus, lower consumption may reflect a signaling effect: since others already know that the individual has received additional income, there is less need for extra consumption to signal wealth or status.

[4] As we explain below, we also include a fourth scenario where respondents receive a bonus with no explicit information about whether coworkers also receive it or whether it is observable. This replicates the typical MPC-style question asked in the literature and is used as a baseline.

[5] The equality $\Delta\bar{c}_{g,-i} = \Delta c_i$ is the consequence of a symmetric equilibrium assumption. Suppose that the reference group is made of only two individuals, i and j (so that $\Delta\bar{c}_{g,-i} = \Delta c_j$). If each individual of the reference group solves the same problem, then $\Delta c_j = \delta\Delta y_j + \alpha\Delta c_i + \phi P_j\Delta y_j$. In the third scenario ($\Delta y_j = \Delta y_i = \Delta y, P_i = P_j = 1$), $\Delta c_j = (\delta + \phi)\Delta y + \alpha\Delta c_i$. It's easy to show that solving the reduced form one arrives at: $\Delta c_j = \Delta c_i = \frac{\delta+\phi}{1-\alpha}\Delta y$.

observability (ϕ); finally, the reduced form parameter β_3 can be used (together with β_1 and β_2) to recover the peer effect parameters α (since $\alpha = \frac{\beta_3}{\beta_1+\beta_2+\beta_3}$) and the implied social multiplier $\frac{1}{1-\alpha}$.

Two distinct channels therefore shape the consumption response to social forces. The first is the observability channel: making the idiosyncratic bonus visible to others shifts the MPC by ϕ, which is empirically signed by the data (through the sign of β_2). The second is the social-interaction channel: a common shock generates a multiplier $\frac{1}{1-\alpha}$, which scales the MPC upwards (as long as $\alpha > 0$). The three scenarios of our survey experiment are designed to recover these parameters *via* randomized assignment.

3. The experimental design

Our analysis uses data from the Italian Survey of Consumer Expectations (ISCE; see Guiso and Jappelli, 2026). ISCE is a large-scale, high-frequency online survey of Italian residents aged 18–75, organized as a quarterly rotating panel. It builds on similar international initiatives, including the Federal Reserve Bank of New York's Survey of Consumer Expectations and the European Central Bank's Consumer Expectations Survey.

The ISCE sample is designed to be representative of the Italian population through stratification by region of residence, age, gender, education, and employment status, following an approach similar to the Bank of Italy's Survey of Household Income and Wealth (SHIW). Participation is incentivized through a combination of individual rewards and charitable donations.[6] The survey achieves a relatively high response rate (around 30–40 percent on average, with some variation across waves) and includes approximately 5,000 respondents per wave.

ISCE has been running since October 2023 and has collected 11 waves of data to date. The panel structure allows tracking individuals over time, with an average retention rate around 80% between consecutive waves. To preserve representativeness, the sample is regularly refreshed on a quarterly basis. Data are collected using Computer Assisted Web Interviewing

[6] The agency designs the incentive scheme to avoid potential self-selection bias, to minimize the influence of incentives on response behavior, and to safeguard data quality. To discourage participation driven solely by financial gain, each personal incentive is paired with a donation to a non-profit charity.

(CAWI). The survey is particularly well suited for our analysis because it incorporates experimental modules designed to study beliefs and behavior.

Appendix C provides further details on the survey and compares the sample means of selected ISCE variables with those from the 2024 SHIW, the most recent available wave. The gender, age, family composition and education levels in both samples are similar (see Table C1), which is noteworthy, given that the ISCE sample includes individuals who are more likely to have internet access and respond to online questionnaires. Consumption medians are generally comparable between the two surveys, but income appears to be underestimated in ISCE relative to SHIW (see Table C2). A possible explanation is that SHIW collects much more detailed information on disposable income than ISCE. Homeownership is reported at 73% in both surveys, while ISCE respondents report higher financial wealth and greater participation in financial markets (bonds, stocks, private pensions and life insurance).

In January 2026 (wave 10 of ISCE) we implement a randomized survey experiment in which respondents are asked to allocate a hypothetical €1,000 bonus across consumption, savings, and transfers over a twelve-month horizon. The bonus is framed as being paid by the employer, which provides a natural, salient and economically meaningful reference group: other employees working in the same workplace. The experiment yields a direct measure of the marginal propensity to consume (MPC), defined as the share of the bonus allocated to spending. The experimental variation maps directly into the model by manipulating two key determinants of consumption: observability and the extent to which the income shock (a firm-level bonus) is shared within the workplace.

All employees are randomly assigned to four scenarios that differ along these dimensions (since the bonus is employer-specific, people who are out of the labor force or self-employed are not part of the research design). In Scenario 1 (baseline), the respondent receives the bonus from the employer with no explicit information about whether coworkers also receive it or whether it is observable. This replicates the typical MPC-style question asked in the literature and can be used as a useful benchmark. In Scenario 2, only the respondent receives the bonus, and he/she is explicitly told that it is not observable by coworkers; this corresponds to the $\{P_i = 0, All_i = 0\}$ scenario described above where people receive a purely idiosyncratic, private shock. In Scenario 3 (the $\{P_i = 1, All_i = 0\}$ case above), the bonus is still idiosyncratic, but it is now publicly observable within the workplace, introducing observability. In Scenario 4 (the

$\{P_i = 1, All_i = 1\}$ case above), all employees receive the same bonus from the employer and it is publicly observable, activating both observability and peer effects.[7]

Note that our identification strategy manipulates the income of coworkers rather than their consumption directly. A more direct approach would be to ask respondents whether observing higher spending by coworkers affects their own consumption. However, questions of this type would make the social context overly salient: respondents who perceive the purpose of the question may adjust their answers accordingly (for instance, overstating the role of peers if they believe social influences are expected, or understating it in order to appear independent-minded) biasing responses through an "experimenter demand effect". In contrast, manipulating whether coworkers receive the same income shock embeds the social channel implicitly (since respondents need to evaluate only their own consumption response to a bonus, with the peer dimension varying across scenarios) and mirrors the structure of actual policy-relevant shocks, such as fiscal transfers that operate on income rather than consumption directly. Nevertheless, this design choice raises three concerns, which we discuss in turn.

First, because the experiment manipulates peer income rather than peer consumption, identification of the social interaction parameter relies on the symmetric equilibrium of the model (as explained in Section 2). In Scenario 4 all coworkers face the same shock, so there exists a unique symmetric equilibrium in which every agent uses the same consumption rule. The mapping from the common income shock to peer consumption is pinned down by the model, and the difference in reported MPC between Scenarios 3 and 4 recovers the parameter α without requiring direct observation of peer spending. This logic would apply equally if we varied the size of the peer bonus (e.g., half or twice the respondent's bonus), which would provide overidentifying restrictions. We chose an equal bonus design because this is the most natural scenario for respondents and because asymmetric bonuses risk confounding the social channel with fairness or inequality aversion issues.

Second, respondents may interpret the shared bonus in Scenario 4 as a signal of firm performance, which could raise MPC through a permanent income mechanism rather than through peer effects (assuming that firm shocks pass through to wages, as shown by Guiso et al. 2005, and many others). However, the survey question explicitly frames the bonus as a hypothetical one-time payment, which should limit interpretation as profit-sharing or as news

[7] We did not include a private-common-shock scenario. The "All receive it" scenario is necessarily public, since when all coworkers receive the same workplace bonus, common knowledge of the shock is the most plausible state.

about future earnings. Moreover, a permanent income revision should raise the allocation to *all* normal goods, including transfers. As we shall see, there is no effect of the common shock variable on the marginal propensity to transfer, which is difficult to reconcile with a permanent income channel.

A final concern is that the shared bonus may trigger a contextual effect of peer income on own consumption (for instance, through relative income concerns) rather than an endogenous effect operating through anticipated peer consumption. This distinction is important, because only the endogenous channel generates a social multiplier effect; a contextual effect of peer income would produce a one-shot level change in consumption but no amplification effects (Manski, 1993). Our survey questions do not allow us to separate the two channels. However, in a consumption setting there is much less scope for a purely contextual channel. In the education literature (where peer effects have been studied widely), the distinction is important because peer background characteristics can affect the outcome (e.g., test scores) production function through mechanisms that are independent of peer achievement. For example, teachers may adjust instruction to the socioeconomic composition of the classroom, high-income parents may lobby for better resources (or donate to improve them), etc., which would act independently of peer test scores. In contrast, in our setting it is much harder to think of theoretical mechanisms that would motivate peers' income as a contextual effect. Relative income concerns, such as feeling poorer because coworkers received a bonus, might appear contextual, but the behavioral response they may trigger (spending more to maintain relative standing) is a response to the anticipated consumption gap, not to peer income *per se*. The same applies to reciprocity norms and informational effects: in each case, peer income matters because it shifts peer consumption, not because it has a direct effect on own consumption.

4. Results

To pave the way for the formal regression analysis below, Figure 1 plots the average MPC in the four groups, and Table 1 reports means, medians and standard deviations of the allocation of the bonus (consumption, transfers, saving) in the four groups. In the baseline and private individual-only conditions (Groups 1 and 2), the MPC is virtually identical (around 0.33), confirming that simply making the social environment explicit (without observability) does not

affect behavior.[8] Introducing observability (Group 3) reduces the MPC to 0.30 and increases saving, consistent with a negative observability effect (social restraint). By contrast, when the bonus is shared and publicly observable (Group 4), the MPC rises to 0.36, suggesting the presence of peer amplification.

The marginal propensity to transfer is very similar in the private and public individual-only scenarios (Groups 2 and 3, 9.5% vs 9.3%, respectively), and declines only slightly to 8.6 percent when the bonus is common knowledge and received by all coworkers (Group 4). This suggests that transfers to family and friends are much less responsive to the treatment conditions than consumption.

The marginal propensity to save closely mirrors the MPC patterns. In particular, saving rises when the bonus is publicly observable but idiosyncratic (Group 3), and declines when the bonus is common knowledge and shared by all coworkers (Group 4). Thus, the fall in saving in the common-shock scenario closely matches the increase in consumption, consistent with the interpretation that peer effects primarily operate through the consumption–saving margin rather than through transfers.

Figure 2 shows the distribution of MPC across respondents and scenarios. The distributions are centered around values between 0.2 and 0.4 and display substantial dispersion, indicating considerable heterogeneity in spending responses across all groups. Figure 3 plots the cumulative distribution function of the MPC for the four groups. The "Public" CDF lies above the others throughout most of the support, reflecting greater mass at low MPCs and suggesting that observability shifts behavior towards consumption restraint in response to income windfalls. In contrast, the "All receive it" CDF lies at the bottom, reflecting greater mass at high MPCs and suggesting the presence of peer amplification effects. Finally, the "No info" and "Only you know" distributions are nearly indistinguishable, consistent with the idea that the standard MPC question is interpreted as a private, idiosyncratic shock. As we shall see, these patterns are confirmed in formal regression analyses.

Table 2 presents balance tests across treatment groups (2,427 employees). A few coefficients are marginally significant for some characteristics (e.g., family size and homeownership), as can occur when multiple balance tests are conducted, but there is no systematic pattern of imbalance and magnitudes are economically negligible. The group-

[8] These average MPC values are broadly consistent with the empirical evidence that find sizable, short-run consumption responses to temporary income shocks, as surveyed by Crawley and Theloudis (2024).

specific joint Wald tests (reported in the last row) rejects the null that observable characteristics jointly predict membership in any treatment group. Finally, an omnibus multinomial test across all four groups yields a p-value of 29.3%. Overall, the results provide no evidence of systematic imbalance across the experimental groups.

We now turn to formal regression analysis. Our estimating regression is equation (5) from above, which we reproduce here:

$$MPC_i = \beta_1 + \beta_2 P_i + \beta_3 All_i + X_i'\gamma + \varepsilon_i \quad (5)$$

where P_i is a dummy equal to one in Scenarios 3 and 4 (observable bonus within the workplace), and All_i is a dummy equal to one in Scenario 4 (bonus paid to all employees). The omitted category is Scenario 2. We also compare Scenarios 1 and 2 to assess how explicitly stating that the bonus is private affects responses relative to standard MPC questions that do not include this qualification. Controls X_i are included for precision but are not required for identification due to random assignment.

Because commonality of the bonus is randomly assigned, the comparison between Scenarios 3 and 4 identifies the causal effect of making a publicly observable income shock common within the workplace. Random assignment rules out bias from correlated individual or workplace characteristics. However, since the experiment manipulates coworkers' income rather than their consumption, this comparison does not, by itself, distinguish an endogenous response to peer consumption from a contextual response to peer income (Manski, 1993). Under the model's symmetry and exclusion assumptions (i.e., that coworkers' receipt of the bonus affects own consumption through its effect on coworkers' consumption), the common shock treatment effect maps into the peer-interaction parameter and the implied social multiplier.

Table 3 reports the regression estimates for different group comparisons. Column 1 compares Scenarios 1 and 2 and shows that the dummy for assignment to Group 2 is close to zero and statistically insignificant. Although this comparison does not bear directly on our experiment, it is useful because Scenario 1 corresponds to the standard MPC question about a transitory income shock: receiving a €1,000 bonus with no further context. Scenario 2, by contrast, makes it explicit that the bonus is private and idiosyncratic. The fact that MPCs are virtually identical in the two scenarios suggests that, when people encounter the standard question without social context, they interpret the shock as private and individual-specific. In

other words, the social environment is not salient in their mental model. Fiscal interventions such as stimulus checks, tax rebates, and broad-based transfers are, however, not atomistic: many people receive them at the same time, and this is typically publicly known. This implies that, to predict the impact of fiscal policy, MPC elicitation should incorporate the relevant social context, as we show below.

Column 2 compares the private, idiosyncratic bonus scenario, Group 2, which serves as the control group, with the public, idiosyncratic bonus scenario, Group 3, thereby isolating the effect of observability. The estimates show that observability has a negative and statistically significant effect of 3.1 percentage points, consistent with norm-driven restraint when receipt of the bonus is publicly known within the workplace.

Column 3 compares Scenario 4 with Scenario 3 and isolates the effect of the bonus being received by all coworkers rather than only by the respondent, holding observability constant. The coefficient on "All receive it" is positive, large, and statistically significant, around 5.6 percentage points. Thus, respondents report higher MPCs when the same bonus is received by all coworkers. This common-shock effect is consistent with peer amplification and, under the maintained assumptions of the model, maps into an endogenous social-interaction effect.

In column 4 we include three groups simultaneously, allowing the separate identification of the observability effect ("Public") and the incremental effect of the common-shock treatment ("All receive it"). This is the estimation of equation (5) without observable controls. The estimates are very stable. Public observability reduces MPC by 3.1 percentage points, while the "All receive it" treatment raises MPC by 5.6 points. These findings align with the model: observability shifts behavior directly, while common workplace shocks activate the social multiplier.

Under the maintained assumptions of the model, the estimates imply a private MPC of $\hat{\delta}$=0.333 (s.e. 0.011), an observability effect of $\hat{\phi}$= −0.031 (s.e. 0.016), and a social-interaction parameter of $\hat{\alpha}$=0.157 (with a robust standard error of 0.041, computed with the delta method), corresponding to a social multiplier of 1.19 (s.e. 0.058). Visibility thus lowers the MPC by about 3 percentage points, while peer interactions amplify the visible MPC by 19 percent. The inclusion of Scenario 1 in column 5 leaves the estimated coefficients and the implied $\hat{\alpha}$ virtually unchanged, reinforcing the view that the standard MPC question is interpreted similarly to the private, idiosyncratic shock scenario.

Table 4 replicates the specifications in Table 3 adding demographic and economic controls. The control variables display intuitive patterns. Men report higher MPCs than women by about 5 percentage points. Middle-aged respondents and larger households report higher MPCs, consistent with greater spending needs. MPCs decline with cash-on-hand, suggesting that households with greater access to liquid resources are less likely to spend the bonus (as found in Jappelli and Pistaferri, 2014). College education, homeownership, and region of residence do not appear to explain differences in MPCs. Most importantly, the main treatment effects are virtually unchanged, indicating that the experimental results are not driven by observable differences in respondent characteristics across scenarios.[9]

One potential concern with the reported MPC measure is the substantial heaping of responses at focal values, as evident from Figure 2. In our data, heaping is particularly pronounced at €500, corresponding to an MPC of 0.5 (22% of the sample reports spending €500 out of the €1,000 bonus). The concentration at €500 is particularly noteworthy, as midpoint responses are often interpreted as reflecting epistemic uncertainty or imprecise beliefs (Bruine de Bruin et al., 2002). Previous research has documented that respondents tend to provide more refined answers at the extremes of a 0–100 scale than in the middle, and that rounding or heaping behavior is systematically related to observable characteristics such as financial sophistication, health, and exposure to macroeconomic events (Giustinelli et al., 2018). Similarly, Gideon et al. (2017) show that responses to quantitative financial questions frequently cluster at round numbers, especially among respondents with lower cognitive ability or when questions are more difficult.

Table B1 in Appendix B shows that the main results are robust to alternative treatments of response quality and focal MPC answers. Column (1) indicates that neither the "Public" nor the "All receive it" treatments significantly predict the probability of reporting the focal response MPC = 0.5. This suggests that treatment effects in the baseline estimates are unlikely to be driven by differential bunching at the midpoint of the response scale. Columns (2) and (3) are directly comparable to the baseline MPC specification (column (4) of Table 4). Excluding respondents who report MPC = 0.5 does not change the main findings: the estimate of the observability effect remains negative and statistically significant (-0.038) and the implied peer

[9] With controls, the intercept is the MPC at $X = 0$ rather than at the sample mean, so we recover the peer effect α using $\frac{\hat{\beta}_3}{(\hat{\beta}_1+\bar{X}'\hat{\gamma})+\hat{\beta}_2+\hat{\beta}_3}$, where $\bar{X}$ is the sample average of the control variables. Standard errors are computed using the delta method from the full covariance matrix of $(\hat{\beta}, \hat{\gamma})$. We assume that covariates shift δ but not α or ϕ.

effect α is positive, highly significant, and similar to the estimate of the baseline model (0.21). The magnitudes are similar when the sample is restricted to respondents who spent at least 10 minutes on the interview. Overall, the table supports the interpretation that the baseline findings are not driven by response heaping or low-quality interviews.

The MPC is bounded between 0 and 1, with mass at both corners. As a robustness check, Table B2 in Appendix B reports two-limit Tobit estimates, with marginal effects computed for the observed MPC. They closely mirror the OLS estimates in Table 3: the observability effect is −3.2 percentage points, the common-shock effect is 5.6 percentage points, and the implied social interaction parameter α is 0.155 (0.041), very close to the estimate of 0.157 in Table 3.

Recall that the survey asks respondents to allocate the hypothetical bonus across three margins: consumption, saving, and transfers to friends and family. Having examined how social interactions shape the MPC, we now turn to the marginal propensity to transfer (MPT), the share of the bonus that respondents would direct towards friends or family members. Table 5 shows that neither observability nor shared workplace shocks have a statistically or economically significant effect on the MPT. The coefficients on both "Public" and "All receive it" are close to zero across all specifications, in contrast with the results for the MPC. At the same time, the average MPT is roughly stable across scenarios (ranging from 8.6% to 9.7%; see Table 1), indicating that transfers represent a non-negligible but treatment-invariant component of how individuals allocate windfall income.

The absence of treatment effects on transfers is consistent with several explanations. Transfers to family and friends are likely driven by altruism, reciprocity, or social obligations, motives for which workplace observability and coworker shocks are largely irrelevant. Relatedly, the relevant reference group for transfers differs from that for consumption: the experimental variation operates within the workplace, while transfers are directed toward family and social contacts outside it.

The fact that the MPT is independent of the social context also speaks to the economic interpretation of the common shock treatment. If respondents interpreted a shared bonus as a signal of firm prosperity, the implied permanent income revision should raise the allocation to *all* normal goods, including transfers. The absence of any treatment effect on transfers is difficult to reconcile with a permanent income channel, and instead supports the interpretation that the common-shock effect on consumption operates through peer interactions.

Overall, the results in this section reveal two opposing social forces operating within the workplace. Observability reduces consumption, consistent with norm-based restraint among coworkers rather than status-seeking behavior. In contrast, shared shocks, when all employees receive the bonus, lead to higher consumption, consistent with peer effects ("keeping-up-with-the-Joneses") within the workplace environment. From a policy perspective, the findings suggest that broad-based transfers within a given economic or organizational environment may generate larger aggregate effects than targeted ones, as shared shocks are amplified through social interactions, whereas idiosyncratic shocks may be dampened by observability.

5. Heterogeneity

In the section above we have estimated a benchmark model with common preference parameters. In practice, the strength of the two key mechanisms identified in the data (observability and peer effects) may differ across workers and workplaces. Some workers may be more susceptible to social pressures or more inclined to "show off"; in some settings the social context may be more salient. To examine these possibilities, in this section we estimate the same specification (5) separately by observable groups:

$$MPC_i = \beta_{1g} + \beta_{2g} P_i + \beta_{3g} All_i + X_i' \gamma_g + \epsilon_i$$

Within group g, β_{2g} measures the visibility effect: the change in the MPC when an idiosyncratic bonus becomes observable to coworkers. The coefficient β_{3g} measures the additional common-shock increment, i.e., the change in the MPC when coworkers also receive the bonus, conditional on the bonus being publicly observable. These estimates capture how the reduced-form coefficients vary across groups.

What these coefficients identify depends on the type of sample split, as discussed in Appendix A. If the split is based on characteristics of workers that define a reference group (such as firm size, which is common to all members of the group) then the estimates can identify group-specific social effects. Under the assumptions spelled out in Appendix A, this allows us to recover firm-size-specific social multipliers and to test, for example, whether peer effects are stronger in smaller firms than in larger firms. The interpretation is different when the split is based on an individual worker characteristic, such as susceptibility to social pressure. Workers who share this trait do not necessarily belong to the same reference group, since they are spread

across different firms. Therefore, one cannot identify a meaningful social multiplier. The exercise is nevertheless informative, because it allows us to test whether the direct response parameters, such as δ_i and ϕ_i, vary systematically with individual traits.

A first dimension of heterogeneity concerns workplace size, which proxies for the intensity of social interactions. If social interactions are stronger in more cohesive or socially proximate environments, peer effects could vary with firm size. In smaller firms, coworkers are more likely to know one another, exchange information about personal circumstances, observe each other's consumption, discuss the size of the bonus, and regard colleagues as a salient reference group. In larger firms, by contrast, employees may interact mainly within narrower teams or departments, weakening the relevance of firm-wide shocks for peer behavior. The visibility effect and the common-shock increment may therefore be expected to be larger in small firms than in large firms. The estimates reported in Table 6 strongly support this hypothesis. The common-shock effect is large and highly significant in firms with fewer than 15 employees (15.3 percentage points), but drops to 4.7 percentage points in firms with 15 to 100 employees and to a negligible 0.8 percentage points in firms with more than 100 employees, with both estimates being statistically insignificant. We reject equality of the "All receive it" coefficients when comparing small versus large firms (p-value 0.7%), and when comparing small versus medium (borderline p-value 8%). The observability effect follows a similar pattern: it is largest in small firms (-10.4 percentage points) and much attenuated in medium and large firms. The implied multipliers are 1.63 (s.e. 0.25) in firms with fewer than 15 employees and 1.03 (s.e. 0.10) in firms with more than 100 (Table 6). This concentration of social effects in small firms is consistent with coworkers being a more salient reference group in smaller workplaces.[10]

A potential concern is that larger social effects in small firms reflect more precise signals about common shocks that workers interpret as shifts in permanent income, rather than consumption peer effects. We argued above that if that were the case, an increase in perceived permanent income would raise both consumption *and* transfers (as long as transfers are normal goods). The average MPT is unaffected by social influences, but this may mask heterogeneity

[10] Information on the number of employees in the firm comes from a previous ISCE wave, conducted in April 2025. As a result, firm size is available only for a smaller sample. Nevertheless, in the subsample for which firm size is observed, the baseline estimates remain very similar to those obtained in the full sample, as shown in column 1 of Table 6. The results are also unchanged when Scenario 1 ("No information") is included in the estimation, confirming that the findings are not sensitive to either the reduced sample or the choice of the comparison group.

by firm size. We therefore apply the same firm-size stratification to the MPT (see Appendix B, Table B3). The coefficient on the "All receive it" variable is small and statistically insignificant in every firm-size bin. A permanent-income reinterpretation of the larger MPC response in small firms would predict a parallel gradient in MPT, but we see none. The firm-size gradient in MPC therefore reflects a channel operating specifically through consumption, consistent with peer interactions and inconsistent with a permanent-income reinterpretation of the shock.

An additional source of heterogeneity comes from differences in social responsiveness. The survey asks respondents to rate on a scale from 1 ("no influence") to 10 ("strong influence") how much an increase in consumption among various social groups would affect their own consumption decisions. We classify as "influenced" those who report any degree of influence from work colleagues (a score above 1) and run separate regressions for the "Not influenced" and "Influenced" groups (columns (1) and (2) of Table 7). The coefficient on "All receive it" is larger for the influenced group (6.6%) than for the not-influenced group (3.6%), though the difference is not statistically significant when controlling for level differences in MPC across groups. The visibility coefficient is negative in both groups but imprecisely estimated.

Overall, these additional results provide complementary evidence on the mechanisms underlying the experimental results and help validate the interpretation of observability and peer effects as important drivers of consumption behavior. Peer effects are not uniform across individuals or environments: they are stronger among socially responsive individuals and in settings where social ties are tighter. Observability exerts a broadly negative effect on consumption that is less dependent on self-reported social sensitivity.

6. Conclusions

This paper studies how consumption responses to transitory income shocks (the MPC) depend on social context, using a randomized survey experiment framed in a workplace setting. Respondents are asked how they would allocate a one-time, unexpected €1,000 bonus received from their employer, with experimental variation in whether the bonus is observable to coworkers and whether it is received individually or by all employees. Under some assumptions, this design allows us to separate two key social mechanisms: observability and peer effects within a well-defined reference group, coworkers.

The results show that these two forces work in opposite directions. Making an idiosyncratic bonus publicly observable within the workplace reduces consumption, consistent

with norm-based restraint or concerns about how spending is perceived by coworkers. By contrast, when the bonus is shared across all employees, consumption increases, reflecting strategic complementarities and a social multiplier: individuals spend more when their coworkers also receive and spend the bonus. In the public, common-shock case, both channels are active, and the overall response reflects the balance between a negative observability effect and a positive peer effect.

Social effects are absent for transfers to family and friends, suggesting that the peer channel operates specifically through consumption. The null effect on transfers also argues against interpreting the common-shock treatment as a signal of firm prosperity (and hence a revision in permanent income), since a permanent income revision should raise the allocation to *all* normal goods, including private transfers. We also find that the common-shock differential is concentrated in firms with fewer than 15 employees, where the implied multiplier is 1.63, against 1.03 in the largest firms.

Our experiment is workplace-specific (a one-time, unexpected employer bonus), and an important question is whether the results would extend to broad fiscal transfers. The common-shock effect we estimate arises in a workplace setting in which coworkers form a salient reference group and the commonality of the shock is explicit. Under our model, this effect maps into a social multiplier operating through coworkers. Government transfers vary in how closely they reproduce these conditions. Relief programs for a specific industry (e.g., manufacturing sectors hit by a "China shock") or a specific occupation (e.g., eliminating taxes on tips received by restaurant and service workers) come close. In a broad tax rebate or a means-tested transfer, the reference group is much larger and the commonality of the shock more uncertain. The social channel is likely attenuated in such cases, so our estimates could be read as an upper bound on peer amplification in fiscal contexts.

Two broader implications follow. First, the standard questions used to elicit MPCs (hypothetical, privately observed, idiosyncratic income changes) may not recover the consumption response to shocks with a different social structure. In our experiment, public observability lowers the response to an idiosyncratic shock, while making a publicly observable shock common to all members of a reference group raises it relative to the public, idiosyncratic case. Second, if the common shock effect operates through peer consumption as posited by the model, consumption models that treat households as atomistic will omit a potentially relevant amplification mechanism.

References

Abel, Andrew B. (1990), "Asset Prices under Habit Formation and Catching Up with the Joneses," *American Economic Review* 80, 38–42.

Becker, Gary S., and Kevin M. Murphy (2000), *Social Economics: Market Behavior in a Social Environment*. Cambridge, MA: Harvard University Press.

Bertrand, Marianne, and Adair Morse (2016), "Trickle-Down Consumption," *Review of Economics and Statistics* 98(5), 863–879.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Rajshri Jayaraman (2015), "Linear Social Interactions Models," *Journal of Political Economy* 123, 444–496.

Bruine de Bruin, Wändi, Paul S. Fischbeck, Neil A. Stiber, and Baruch Fischhoff (2002), "What Number is 'Fifty-Fifty'? Redistributing Excessive 50% Responses in Elicited Probabilities," *Risk Analysis* 22, 713–723.

Bursztyn, Leonardo, Florian Ederer, Bruno Ferman, and Noam Yuchtman (2014), "Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions," *Econometrica* 82(4), 1273–1301.

Bursztyn, Leonardo, Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao (2018), "Status Goods: Experimental Evidence from Platinum Credit Cards," *Quarterly Journal of Economics* 133(3), 1561–1595.

Bursztyn, Leonardo, and Robert Jensen (2015), "How Does Peer Pressure Affect Educational Investments?" *Quarterly Journal of Economics* 130(3), 1329–1367.

Carrell, Scott E., Richard L. Fullerton, and James E. West (2009), "Does Your Cohort Matter? Measuring Peer Effects in College Achievement," *Journal of Labor Economics* 27(3), 439–464.

Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov (2009), "Conspicuous Consumption and Race," *Quarterly Journal of Economics* 124(2), 425–467.

Crawley, Edmund, and Alexandros Theloudis (2024), "Income Shocks and Their Transmission into Consumption," *Finance and Economics Discussion Series* No. 2024-038.

De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri (2020), "Consumption Network Effects," *Review of Economic Studies* 87(1), 130–163.

Duesenberry, James S. (1949), *Income, Saving and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.

Fuster, Andreas, Greg Kaplan, and Basit Zafar (2021), "What Would You Do with a $500 Unexpected Gain? Spending Responses and Marginal Propensities to Consume," *American Economic Review: Insights* 3(4), 647–665.

Gideon, Michael, Brooke Helppie-McFall, and Joanne W. Hsu (2017), "Heaping at Round Numbers on Financial Questions: The Role of Satisficing," *Survey Research Methods* 11, 189–214.

Giustinelli, Pamela, Charles F. Manski, and Francesca Molinari (2018), "Tail and Center Rounding of Probabilistic Expectations in the Health and Retirement Study," NBER Working Paper No. 24559.

Glaeser, Edward L., Bruce Sacerdote, and José A. Scheinkman (2003), "The Social Multiplier," *Journal of the European Economic Association* 1(2–3), 345–353.

Guiso, Luigi, and Tullio Jappelli (2026), "The Italian Survey of Consumer Expectations: Statistical Bulletin," CSEF Working Paper No. 722.

Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi (2005), "Insurance within the Firm," *Journal of Political Economy* 113(5), 1054–1087.

Jappelli, Tullio, and Luigi Pistaferri (2014), "Fiscal Policy and MPC Heterogeneity," *American Economic Journal: Macroeconomics* 6(4), 107–136.

Jappelli, Tullio, and Luigi Pistaferri (2020), "The Consumption Response to Income Changes," *Annual Review of Economics* 12, 479–506.

Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn (2011), "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery," *American Economic Review* 101(5), 2226–2247.

Manski, Charles F. (1993), "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* 60(3), 531–542.

Sacerdote, Bruce (2001), "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *Quarterly Journal of Economics* 116(2), 681–704.

Stantcheva, Stefanie (2023), "How to Run Surveys: A Guide to Creating Your Own Identifying Variation and Revealing the Invisible," *Annual Review of Economics* 15(1), 205-234.

Veblen, Thorstein (1899), *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions*. New York: Macmillan.

Figure 1. Average MPC by treatment groups

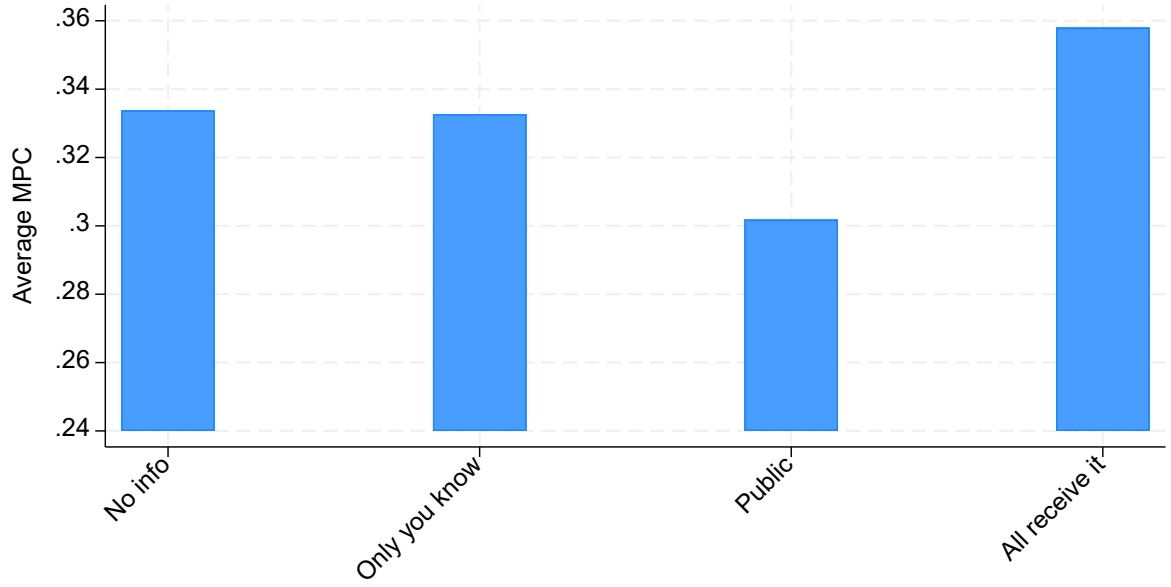

Notes: The figure reports the average marginal propensity to consume (MPC) in the control and treatments groups. The MPC is defined as the share of a hypothetical €1,000 bonus allocated to consumption over the next 12 months.

Figure 2. Distribution of MPC by treatment groups

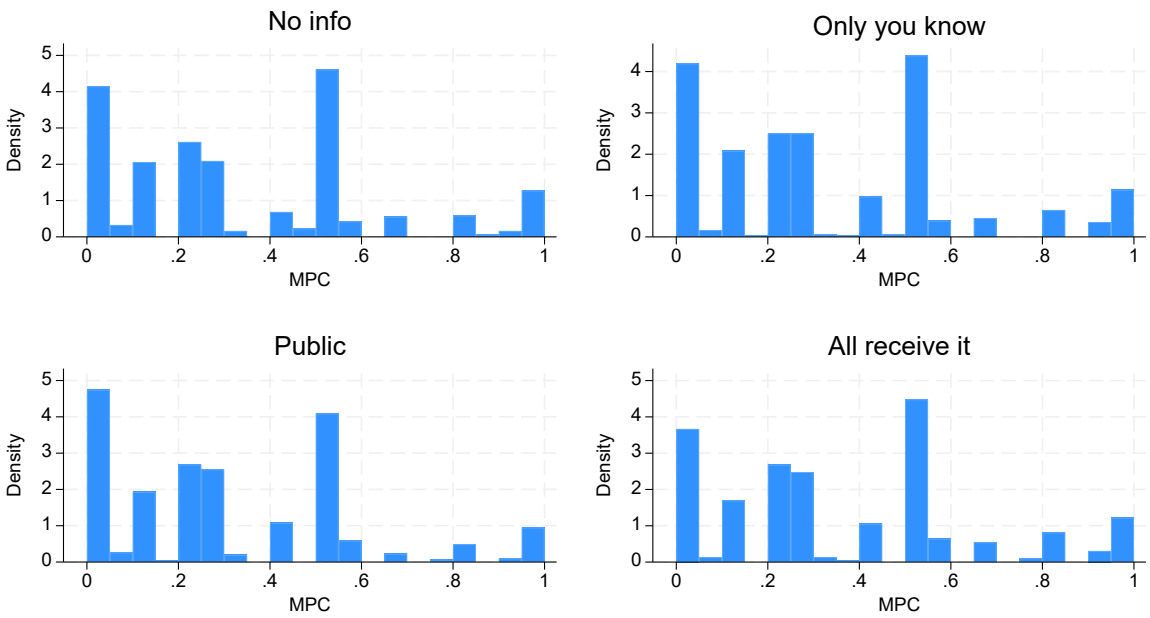

Notes: The figure reports, for each group, the distribution of the marginal propensity to consume (MPC), defined as the share of a hypothetical €1,000 bonus allocated to consumption over the next 12 months.

Figure 3. The CDF of the MPC by treatment groups

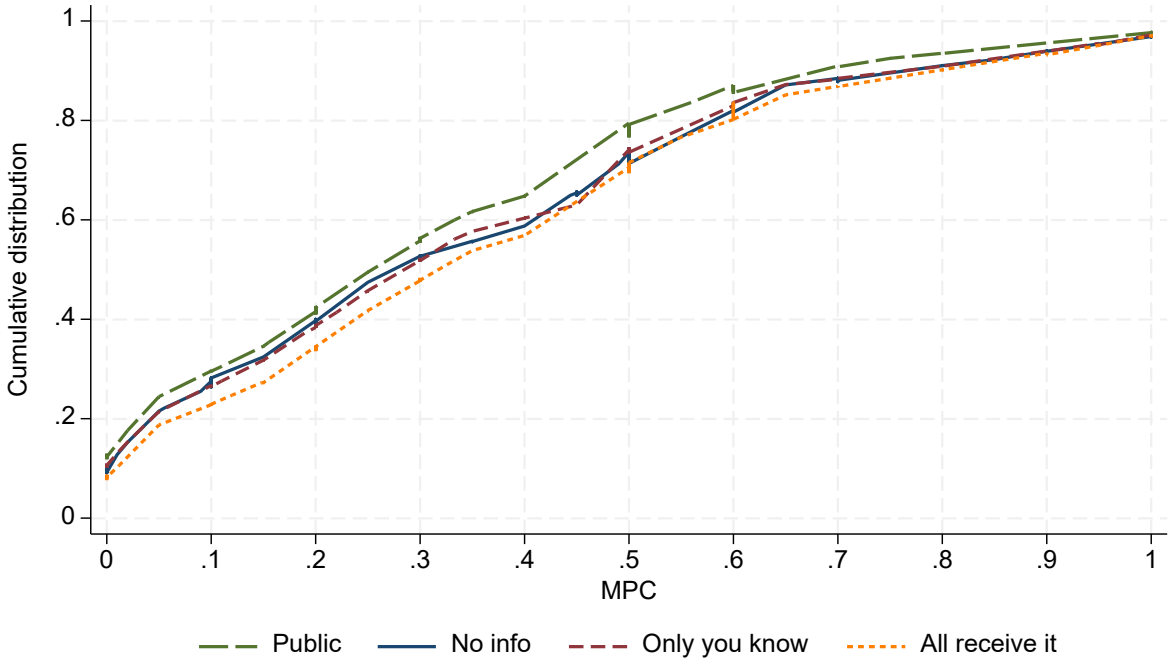

Notes: Cumulative distribution of the marginal propensity to consume (MPC) by treatment scenario. The €1,000 bonus is allocated to consumption over 12 months.

Table 1. Sample statistics

	Mean	Median	Standard deviation
Group 1 - No additional information			
Marginal propensity to consume	0.334	0.3	0.287
Marginal propensity to save	0.569	0.5	0.309
Marginal propensity to transfer	0.097	0	0.175
Group 2 - Given only to you and only you know it			
Marginal propensity to consume	0.333	0.3	0.285
Marginal propensity to save	0.572	0.5	0.308
Marginal propensity to transfer	0.095	0	0.173
Group 3 - Given only to you and everyone knows it			
Marginal propensity to consume	0.302	0.25	0.269
Marginal propensity to save	0.606	0.6	0.302
Marginal propensity to transfer	0.093	0	0.158
Group 4 - Given to everyone and everyone knows			
Marginal propensity to consume	0.358	0.3	0.286
Marginal propensity to save	0.556	0.5	0.311
Marginal propensity to transfer	0.086	0	0.157

Notes: This table reports summary statistics for the marginal propensities to consume, save, and transfer out of a hypothetical €1,000 bonus, by experimental group. Respondents are randomly assigned to four scenarios: (1) no additional information; (2) bonus received only by the respondent and privately observed; (3) bonus received only by the respondent but publicly observable; (4) bonus received by all coworkers and publicly observable. The marginal propensities are defined as the share of the €1,000 bonus allocated to each use over a 12-month horizon.

Table 2. Balance test

	No info	Only you know	Public	All receive it
Male	0.013	-0.007	-0.015	0.009
	(0.018)	(0.018)	(0.018)	(0.018)
Age 35 to 50	-0.018	-0.031	0.055	-0.005
	(0.024)	(0.023)	(0.024)**	(0.023)
Age 51 to 65	0.007	-0.028	0.020	-0.000
	(0.025)	(0.024)	(0.025)	(0.024)
Age 66 to 75	0.145	-0.171	-0.027	0.011
	(0.090)	(0.110)	(0.102)	(0.097)
Family size	0.004	-0.017	0.004	0.009
	(0.008)	(0.008)**	(0.008)	(0.008)
College degree	-0.031	-0.018	0.031	0.015
	(0.020)	(0.020)	(0.020)	(0.019)
North	-0.001	-0.026	0.008	0.019
	(0.021)	(0.021)	(0.021)	(0.021)
Centre	0.002	-0.012	-0.025	0.034
	(0.025)	(0.025)	(0.025)	(0.024)
Log cash-on-hand	0.001	0.005	-0.004	-0.002
	(0.008)	(0.008)	(0.008)	(0.008)
Homeowner	-0.022	-0.044	0.023	0.046
	(0.022)	(0.022)**	(0.023)	(0.023)**
N	2,427	2,427	2,427	2,427
Joint p-value	0.501	0.146	0.246	0.565

Notes: This table reports regressions of the probability that respondents are included in one of the four groups of the experiment: (1) no additional information; (2) bonus received only by the respondent and privately observed; (3) bonus received only by the respondent but publicly observable; (4) bonus received by all coworkers and publicly observable. Robust standard errors are reported in parentheses. The last row reports the joint Wald p-values from each probit, testing the null that demographic characteristics do not jointly predict assignment to the corresponding group. *,**,***: significant at 10%, 5% and 1%, respectively.

Table 3. Regressions for the MPC

	(1)	(2)	(3)	(4)	(5)
	Group 2 vs 1	Group 3 vs 2	Groups 4 vs 3	Groups 2-3-4	All groups
Only you know	-0.001				-0.001
	(0.016)				(0.016)
Public		-0.031		-0.031	-0.032
		(0.016)*		(0.016)*	(0.016)**
All receive it			0.056	0.056	0.056
			(0.016)***	(0.016)***	(0.016)***
Constant	0.334	0.333	0.302	0.333	0.334
	(0.012)***	(0.011)***	(0.011)***	(0.011)***	(0.012)***
Implied α				0.157	0.157
				(0.041)***	(0.041)***
Social multiplier				1.186	1.186
				(0.058)	(0.058)
R^2	0.00	0.00	0.01	0.01	0.00
N	1,227	1,226	1,200	1,815	2,427

Notes: This table reports OLS estimates of the marginal propensity to consume (MPC), defined as the share of a €1,000 bonus allocated to consumption. Respondents are randomly assigned to four scenarios: (1) no additional information; (2) bonus received only by the respondent and privately observed; (3) bonus received only by the respondent but publicly observable; (4) bonus received by all coworkers and publicly observable. "Only you know" is a dummy equal to one for the private, idiosyncratic bonus scenario (Scenario 2). "Public" is a dummy equal to one when the bonus is observable to coworkers (Scenarios 3 and 4), and "All receive it" is a dummy equal to one when the bonus is received by all coworkers (Scenario 4). Column 1 uses observations from Scenarios 1 and 2, with Scenario 1 as the omitted category. Column 2 uses observations from Scenarios 2 and 3, with Scenario 2 as the omitted category. Column 3 uses observations from Scenarios 3 and 4, with Scenario 3 as the omitted category. Column 4 uses observations from Scenarios 2, 3, and 4, with Scenario 2 as the omitted category. Column 5 uses observations from Scenarios 1, 2, 3, and 4, with Scenario 1 as the omitted category. Robust standard errors are reported in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Table 4. Regressions for the MPC, with demographics

	Group 2 vs. 1	Group 3 vs 2	Group 4 vs 3	Groups 2-3-4	All groups
Only you know	0.003				0.001
	(0.016)				(0.016)
Public		-0.032		-0.033	-0.031
		(0.016)**		(0.016)**	(0.016)**
All receive it			0.058	0.057	0.057
			(0.016)***	(0.016)***	(0.016)***
Male	0.044	0.056	0.049	0.051	0.047
	(0.016)***	(0.016)***	(0.016)***	(0.013)***	(0.011)***
Age 35 to 50	0.041	0.031	0.057	0.039	0.050
	(0.020)**	(0.020)	(0.019)***	(0.016)**	(0.014)***
Age 51 to 65	0.062	0.023	0.057	0.043	0.061
	(0.021)***	(0.021)	(0.021)***	(0.017)**	(0.015)***
Age 66 to 75	0.134	0.092	-0.123	-0.084	0.024
	(0.065)**	(0.145)	(0.117)	(0.105)	(0.068)
Family size	0.015	0.017	0.011	0.015	0.013
	(0.008)*	(0.008)**	(0.008)	(0.006)**	(0.005)**
College degree	0.009	-0.021	0.018	0.010	0.013
	(0.018)	(0.018)	(0.018)	(0.015)	(0.013)
North	-0.042	-0.010	0.000	-0.008	-0.022
	(0.019)**	(0.018)	(0.018)	(0.015)	(0.013)
Centre	-0.009	-0.002	-0.016	-0.012	-0.013
	(0.023)	(0.023)	(0.022)	(0.018)	(0.016)
Log cash-on-hand	-0.014	-0.012	-0.028	-0.024	-0.020
	(0.008)*	(0.008)	(0.008)***	(0.007)***	(0.006)***
Homeowner	-0.028	-0.030	-0.015	-0.024	-0.022
	(0.020)	(0.021)	(0.023)	(0.018)	(0.015)
Constant	0.397	0.380	0.482	0.486	0.450
	(0.081)***	(0.081)***	(0.077)***	(0.065)***	(0.056)***
Implied α				0.160	0.159
				(0.041)***	(0.041)***
Social multiplier				1.190	1.189
				(0.058)	(0.058)
R^2	0.03	0.03	0.04	0.03	0.03
N	1,227	1,226	1,200	1,815	2,427

Notes: This table reports OLS estimates of the marginal propensity to consume (MPC), defined as the share of a €1,000 bonus allocated to consumption. Respondents are randomly assigned to four scenarios: (1) no additional information; (2) bonus received only by the respondent and privately observed; (3) bonus received only by the respondent but publicly observable; (4) bonus received by all coworkers and publicly observable. "Only you know" is a dummy equal to one for the private, idiosyncratic bonus scenario (Scenario 2). "Public" is a dummy equal to one when the bonus is observable to coworkers (Scenarios 3 and 4), and "All receive it" is a dummy equal to one when the bonus is received by all coworkers (Scenario 4). Column 1 uses observations from Scenarios 1 and 2, with Scenario 1 as the omitted category. Column 2 uses observations from Scenarios 2 and 3, with Scenario 2 as the omitted category. Column 3 uses observations from Scenarios 3 and 4, with Scenario 3 as the omitted category. Column 4 uses observations from Scenarios 2, 3, and 4, with Scenario 2 as the omitted category. Column 5 uses observations from Scenarios 1, 2, 3, and 4, with Scenario 1 as the omitted category. Robust standard errors are reported in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Table 5. Regressions for the Marginal Propensity to Transfer (MPT)

	Group 2 vs. 1	Group 3 vs. 2	Group 4 vs. 3	Groups 2-3-4	All groups
Only you know	0.001				-0.000
	(0.010)				(0.010)
Public		-0.004		-0.003	-0.004
		(0.009)		(0.009)	(0.009)
All receive it			-0.006	-0.006	-0.006
			(0.009)	(0.009)	(0.009)
Male	0.000	-0.016	-0.018	-0.012	-0.009
	(0.010)	(0.010)	(0.009)**	(0.008)	(0.007)
Age 35 to 50	0.019	0.013	-0.003	0.004	0.008
	(0.012)	(0.012)	(0.011)	(0.009)	(0.008)
Age 51 to 65	0.048	0.035	0.013	0.028	0.031
	(0.013)***	(0.013)***	(0.013)	(0.011)***	(0.009)***
Age 66 to 75	0.072	0.009	-0.032	-0.026	0.027
	(0.042)*	(0.036)	(0.035)	(0.029)	(0.029)
Family size	0.019	0.017	0.014	0.015	0.016
	(0.004)***	(0.004)***	(0.004)***	(0.003)***	(0.003)***
College degree	0.026	0.006	-0.021	-0.002	0.003
	(0.011)**	(0.010)	(0.009)**	(0.008)	(0.007)
North	-0.008	-0.031	-0.036	-0.031	-0.022
	(0.012)	(0.011)***	(0.011)***	(0.009)***	(0.008)***
Centre	-0.012	-0.030	-0.027	-0.023	-0.020
	(0.014)	(0.013)**	(0.013)**	(0.011)**	(0.010)**
Log cash-on-hand	-0.004	0.002	-0.003	0.000	-0.004
	(0.005)	(0.005)	(0.004)	(0.004)	(0.003)
Homeowner	-0.035	-0.031	-0.010	-0.023	-0.024
	(0.014)**	(0.013)**	(0.012)	(0.011)**	(0.010)**
Constant	0.081	0.067	0.129	0.083	0.108
	(0.047)*	(0.044)	(0.045)***	(0.036)**	(0.033)***
R^2	0.03	0.03	0.03	0.03	0.03
N	1,227	1,226	1,200	1,815	2,427

Notes: This table reports OLS estimates of the marginal propensity to transfer, defined as the share of a €1,000 bonus allocated to relatives or friends. Respondents are randomly assigned to four scenarios: (1) no additional information; (2) bonus received only by the respondent and privately observed; (3) bonus received only by the respondent but publicly observable; (4) bonus received by all coworkers and publicly observable. "Only you know" is a dummy equal to one for the private, idiosyncratic bonus scenario (Scenario 2). "Public" is a dummy equal to one when the bonus is observable to coworkers (Scenarios 3 and 4), and "All receive it" is a dummy equal to one when the bonus is received by all coworkers (Scenario 4). Column 1 uses observations from Scenarios 1 and 2, with Scenario 1 as the omitted category. Column 2 uses observations from Scenarios 2 and 3, with Scenario 2 as the omitted category. Column 3 uses observations from Scenarios 3 and 4, with Scenario 3 as the omitted category. Column 4 uses observations from Scenarios 2, 3, and 4, with Scenario 2 as the omitted category. Column 5 uses observations from Scenarios 1, 2, 3, and 4, with Scenario 1 as the omitted category. Robust standard errors are reported in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Table 6. MPC heterogeneity: Firm size

	(1) All firm sizes	(2) <15 employees	(3) 15-100 employees	(4) >100 employees
Public	-0.038	-0.104	-0.029	-0.022
	(0.021)*	(0.046)**	(0.045)	(0.028)
All receive it	0.053	0.153	0.047	0.008
	(0.021)**	(0.045)***	(0.041)	(0.030)
Male	0.063	0.016	0.045	0.095
	(0.018)***	(0.038)	(0.036)	(0.024)***
Age 35 to 50	0.054	0.000	0.033	0.087
	(0.021)**	(0.048)	(0.041)	(0.030)***
Age 51 to 65	0.045	-0.003	0.027	0.072
	(0.024)*	(0.056)	(0.046)	(0.033)**
Age 66 to 75	-0.104	-0.321	0.033	-0.253
	(0.129)	(0.108)***	(0.207)	(0.045)***
Family size	0.019	0.010	0.008	0.025
	(0.008)**	(0.021)	(0.015)	(0.011)**
College degree	-0.003	-0.032	-0.005	0.007
	(0.021)	(0.050)	(0.038)	(0.029)
North	0.020	0.038	0.045	0.002
	(0.021)	(0.045)	(0.038)	(0.029)
Centre	0.027	0.041	0.034	0.023
	(0.024)	(0.056)	(0.048)	(0.033)
Log cash-on-hand	-0.027	-0.023	-0.029	-0.020
	(0.009)***	(0.021)	(0.018)	(0.012)*
Homeowner	-0.011	-0.070	-0.006	0.017
	(0.023)	(0.052)	(0.049)	(0.031)
Constant	0.463	0.569	0.549	0.309
	(0.091)***	(0.216)***	(0.171)***	(0.124)**
Implied social multiplier	1.177	1.631	1.140	1.028
	(0.077)	(0.245)	(0.134)	(0.102)
R^2	0.05	0.09	0.03	0.07
N	987	216	271	500

Note. This table reports heterogeneous treatment effects of observability ("Public") and shared shocks ("All receive it") on the MPC. Columns split the sample based on firm size (number of employees). Information on number of employees in the firm is taken from a previous wave of ISCE (April 2025), and therefore the sample drops. Robust standard errors are in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Table 7. MPC heterogeneity: Social influence

	(1) Not influenced	(2) Influenced
Public	-0.050	-0.023
	(0.029)*	(0.019)
All receive it	0.036	0.066
	(0.028)	(0.019)***
Male	0.075	0.034
	(0.024)***	(0.016)**
Age 35 to 50	0.047	0.040
	(0.028)*	(0.019)**
Age 51 to 65	0.090	0.027
	(0.030)***	(0.021)
Age 66 to 75	0.036	-0.261
	(0.133)	(0.088)***
Family size	0.020	0.010
	(0.011)*	(0.007)
College degree	0.044	-0.009
	(0.027)*	(0.018)
North	0.033	-0.018
	(0.027)	(0.018)
Centre	0.007	-0.014
	(0.034)	(0.022)
Log cash-on-hand	-0.030	-0.021
	(0.012)**	(0.008)***
Homeowner	-0.023	-0.020
	(0.032)	(0.021)
Constant	0.435	0.504
	(0.116)***	(0.079)***
R^2	0.06	0.03
N	607	1,208

Note. This table reports heterogeneous treatment effects of observability ("Public") and shared shocks ("All receive it") on the MPC. Columns split the sample based on self-reported susceptibility to social influence. "Influenced" is a dummy variable equal to one if the respondents report to be influenced by coworkers in consumption decisions. Robust standard errors are in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Appendix A: Heterogeneity

This Appendix derives what the experimental design identifies when the structural parameters vary across individuals. Suppose that the parameters of the linearized response of Section 2 are heterogeneous across individuals:

$$\Delta c_i = (\delta_i + \phi_i P_i)\,\Delta y_i + \alpha_i\,\Delta \bar{c}_{g,-i} \tag{B.1}$$

where $0 \leq \alpha_i < 1$ for all i. The bonus is the fixed amount Δy in every scenario, so $\text{MPC}_i = \Delta c_i / \Delta y$.

In the private idiosyncratic scenario (where $P_i = 0$, $\Delta \bar{c}_{g,-i} = 0$), $\text{MPC}_i = \delta_i$. In the public idiosyncratic scenario ($P_i = 1, \Delta \bar{c}_{g,-i} = 0$), $\text{MPC}_i = \delta_i + \phi_i$. Finally, in the public common scenario, every member of i's reference group g (a group of size n_g) receives Δy and $P_i = 1$, so (B.1) can be rewritten as:

$$\Delta c_i = (\delta_i + \phi_i)\,\Delta y + \alpha_i\,\Delta \bar{c}_g \tag{B.2}$$

where $\Delta \bar{c}_g$ is the mean consumption change of all coworkers.[11] Averaging (B.2) over the members of g then gives $\Delta \bar{c}_g = \left(\bar{\delta}_g + \bar{\phi}_g\right)\Delta y + \bar{\alpha}_g\,\Delta \bar{c}_g$, and hence,

$$\Delta \bar{c}_g = \frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\,\Delta y \tag{B.3}$$

Individual MPC's can be obtained substituting (B.3) into (B.2) to yield:

$$\text{MPC}_i = \delta_i + \phi_i + \alpha_i\,\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}$$

In our experiments, scenarios are randomly assigned, independently of $(\delta_i, \phi_i, \alpha_i)$ and of group composition. The population average MPC's in the three scenarios are, respectively:

$$\mu_{00} = E(MPC_i | P_i = 0, All_i = 0) = E(\delta_i) \tag{B.4.1}$$

$$\mu_{10} = E(MPC_i | P_i = 1, All_i = 0) = E(\delta_i) + E(\phi_i) \tag{B.4.2}$$

$$\mu_{11} = E(MPC_i | P_i = 1, All_i = 1) = E(\delta_i) + E(\phi_i) + E\left[\alpha_i \frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\right] \tag{B.4.3}$$

---

[11] Strictly speaking, i responds to the leave-one-out mean $\Delta \bar{c}_{g,-i} = \Delta \bar{c}_g + \frac{\Delta \bar{c}_g - \Delta c_i}{n_g - 1}$. Substitute this expression in (B.1) and consider the case $P = 1$ and $\Delta y_i = \Delta y$, so that $\Delta c_i = (\delta_i + \phi_i)\,\Delta y + \alpha_i\,\Delta \bar{c}_g + \alpha_i \frac{\Delta \bar{c}_g - \Delta c_i}{n_g - 1}$. Take now the average over all members of the reference group to yield: $\Delta \bar{c}_g = \left(\bar{\delta}_g + \bar{\phi}_g\right)\Delta y + \bar{\alpha}_g\,\Delta \bar{c}_g + \frac{1}{n_g - 1}\left(\bar{\alpha}_g \Delta \bar{c}_g - \frac{\sum_i \alpha_i\,\Delta c_i}{n_g}\right)$. The last term is (minus) the covariance between α_i and Δc_i, scaled by $(n_g - 1)$. If α_i is constant within the firm, the correction term is exactly 0 and can be ignored throughout.

Now we use the law of iterated expectations ($E(x_i) = E(\bar{x}_g)$).[12] So:

$$\mu_{00} = E(MPC_i|P_i = 0, All_i = 0) = E(\bar{\delta}_g)$$

$$\mu_{10} = E(MPC_i|P_i = 1, All_i = 0) = E(\bar{\delta}_g) + E(\bar{\phi}_g)$$

$$\mu_{11} = E(MPC_i|P_i = 1, All_i = 1) = E(\bar{\delta}_g) + E(\bar{\phi}_g) + E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\bar{\alpha}_g\right) = E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\right)$$

Because of random assignment, a reduced form regression of MPC_i on a constant term, P_i and All_i is:

$$MPC_i = \beta_1 + \beta_2 P_i + \beta_3 All_i = \mu_{00} + (\mu_{10} - \mu_{00})P_i + (\mu_{11} - \mu_{10})All_i \quad \text{(B.5)}$$

Hence, β_1 recovers $\mu_{00} = E(\bar{\delta}_g) = E(\delta_i)$, β_2 recovers $\mu_{10} - \mu_{00} = E(\bar{\phi}_g) = E(\phi_i)$, and β_3 recovers $E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\right) - E(\bar{\delta}_g + \bar{\phi}_g) = E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g} - (\bar{\delta}_g + \bar{\phi}_g)\right) = E\left(\frac{\bar{\alpha}_g}{1 - \bar{\alpha}_g}(\bar{\delta}_g + \bar{\phi}_g)\right)$.

A key question is whether it is possible to identify a summary measure of the social multiplier. It is easy to show that:

$$\frac{\mu_{11}}{\mu_{10}} = \frac{E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{1 - \bar{\alpha}_g}\right)}{E(\bar{\delta}_g + \bar{\phi}_g)} = E\left(\frac{\bar{\delta}_g + \bar{\phi}_g}{E(\bar{\delta}_g + \bar{\phi}_g)}\frac{1}{1 - \bar{\alpha}_g}\right) = E\left(\omega_g \frac{1}{1 - \bar{\alpha}_g}\right),$$

Hence, in our case the ratio of the average MPC for the $\{P = 1, All = 1\}$ group and the average MPC for the $\{P = 1, All = 0\}$ group identifies a *weighted* average of the group-specific multipliers (where the weights are given by the group's responsiveness to the bonus -the regular MPC- plus the group-specific response to the publicness of the bonus (normalized by its population average so that the weights are nonnegative and average to one)). If $(\bar{\delta}_g + \bar{\phi}_g)$ is mean independent of $\bar{\alpha}_g$ across groups, the weights drop out and we recover the average social multiplier.

In general, the pooled estimate identifies a weighted average of group-specific social multipliers, with greater weight placed on groups that have higher baseline MPCs (because ω_g is increasing in δ_g, assuming $\phi_g$is fixed, or not negatively correlated enough to offset δ_g).

[12] The empirical analog is $n^{-1}\sum_{i=1}^{n} x_i = \sum_g \left(\frac{n_g}{n}\right)\left(n_g^{-1}\sum_{i \in g} x_i\right)$, i.e., the population mean of x_i is the average of the group means weighted by the relative size of each group in the population, or $E(x_i) = E(\bar{x}_g)$.

Therefore, the weighted average needs not coincide with the simple average of the group-specific multipliers.

Sample splits. The identification above is preserved when the split is done with respect to characteristics that define peers. In our case, peers are co-workers, i.e., people who work for the same firm. Hence, when we split the sample by firm size, the result above holds "cell by cell" (where each cell is a firm size grouping). This means that each cell X identifies its own weighted multiplier $\frac{\mu_{11}(X)}{\mu_{10}(X)}$, and the pooled multiplier is a weighted average of the cell multipliers.

In contrast, if the split variable is a characteristic of the worker (such as self-reported susceptibility to peer effects, say Z), the conditioning selects workers, not workplaces. The conditional cell means $\mu_{00}(Z)$ and $\mu_{10}(Z)$ identify $E(\delta_i|Z)$ and $E(\delta_i + \varphi_i|\, Z)$, but the common-shock differential is $\mu_{11}(Z) - \mu_{10}(Z) = E[\alpha_i \frac{\bar{\delta}_g + \bar{\varphi}_g}{1 - \bar{\alpha}_g}|Z]$. We can no longer recover a meaningful social multiplier because the cell ratios mix individual susceptibility with the equilibrium response of the reference group.

To see this with a simple numerical example, suppose that in each firm half of workers are susceptible to social pressures from coworkers ($\alpha_i = 0.4$) and half are not ($\alpha_i = 0$). Assume also that there is no other heterogeneity; in particular, $\delta_i = 0.4$ and $\phi_i = 0$ for all workers. This means that $\bar{\alpha}_g = 0.2$ in each firm, there is a common social multiplier of $\frac{1}{1-\bar{\alpha}_g} = 1.25$, and each firm's response to the common bonus is $\frac{\bar{\delta}_g + \bar{\phi}_g}{1-\bar{\alpha}_g} = 0.5$. Consider now running regressions splitting the sample between susceptibles and non-susceptibles to peer pressure. Using (B.4.1)-(B.4.3), among the susceptibles $\mu_{00} = \mu_{10} = 0.4$ and $\mu_{11} = 0.4 + 0.4 \times 0.5 = 0.6$; while among the non-susceptibles $\mu_{00} = \mu_{10} = 0.4$ and $\mu_{11} = 0.4 + 0 \times 0.5 = 0.4$. The implied social multiplier μ_{11}/μ_{10} is therefore 1.5 in the susceptible group and 1.0 in the non-susceptible group. Neither is the multiplier of any reference group (co-worker) in the economy. This is because the cell ratios track individual susceptibility to peer pressure, not group equilibrium.

Online Appendix B: Additional Tables

Table B1. Robustness checks

	Prob(MPC=0.5)	Drop MPC=0.50	Time>=10min
	(1)	(2)	(3)
Public	-0.017	-0.038	-0.034
	(0.023)	(0.019)**	(0.018)*
All receive it	0.021	0.068	0.063
	(0.024)	(0.019)***	(0.019)***
Male	0.015	0.060	0.053
	(0.020)	(0.016)***	(0.015)***
Age 35 to 50	0.008	0.046	0.035
	(0.026)	(0.019)**	(0.019)*
Age 51 to 65	-0.001	0.055	0.052
	(0.027)	(0.021)***	(0.020)***
Age 66 to 75	-0.213	-0.033	-0.076
	(0.025)***	(0.105)	(0.115)
Family size	0.012	0.016	0.017
	(0.009)	(0.007)**	(0.007)**
College degree	-0.013	0.020	0.016
	(0.022)	(0.018)	(0.017)
North	-0.009	-0.009	-0.005
	(0.023)	(0.018)	(0.017)
Centre	-0.048	-0.005	-0.009
	(0.027)*	(0.022)	(0.021)
Log cash-on-hand	-0.006	-0.030	-0.022
	(0.009)	(0.008)***	(0.007)***
Homeowner	0.014	-0.031	-0.015
	(0.025)	(0.022)	(0.020)
Constant	0.246	0.484	0.444
	(0.088)***	(0.079)***	(0.073)***
Implied α		0.214	0.175
		(0.072)***	(0.060)***
Social multiplier		1.273	1.212
		(0.087)	(0.069)
R^2	0.01	0.04	0.04
N	1,815	1,423	1,373

Notes: This table reports robustness checks for the baseline OLS specification of the marginal propensity to consume (MPC). Column (1) reports estimates of the probability of reporting MPC=0.5. Column (2) re-estimates the main regression of column (4), Table 3, excluding respondents with focal responses at MPC = 0.5. Column (3) restricts the sample to respondents with interview duration greater than 10 minutes, used as a proxy for higher-quality survey responses. "Public" is a dummy equal to one when the bonus is observable to coworkers (Scenarios 3 and 4), and "All receive it" is a dummy equal to one when the bonus is received by all coworkers (Scenario 4). All specifications include the same set of demographic controls as in Table 4, column (4). The omitted category is the private individual-only scenario (Scenario 2). Robust standard errors are reported in parentheses. *,**,***: significant at 10%, 5% and 1%, respectively.

Table B2. Tobit regressions for the MPC

	(1) Groups 2 vs. 1	(2) Groups 3 vs. 2	(3) Groups 4 vs. 3	(4) Groups 2-3-4	(5) All groups
Only you know	0.0017				
	(0.0164)				
Public		-0.0313**		-0.0322**	-0.0324**
		(0.0159)		(0.0161)	(0.0140)
All receive it			0.0569***	0.0562***	0.0563***
			(0.0161)	(0.0161)	(0.0162)
Male	0.0440***	0.0544***	0.0457***	0.0486***	0.0454***
	(0.0166)	(0.0161)	(0.0163)	(0.0134)	(0.0117)
Age 35 to 50	0.0430**	0.0238	0.0472**	0.0316*	0.0465***
	(0.0203)	(0.0200)	(0.0197)	(0.0163)	(0.0142)
Age 51 to 65	0.0609***	0.0179	0.0467**	0.0346*	0.0556***
	(0.0219)	(0.0214)	(0.0218)	(0.0178)	(0.0155)
Age 66 to 75	0.122**	0.0724	-0.207	-0.148	-0.00164
	(0.0583)	(0.148)	(0.167)	(0.133)	(0.0742)
Family size	0.0158**	0.0178**	0.0123	0.0161**	0.0141***
	(0.00772)	(0.00749)	(0.00772)	(0.00626)	(0.00546)
College degree	0.0131	-0.0230	0.0128	0.00722	0.0119
	(0.0183)	(0.0178)	(0.0184)	(0.0149)	(0.0130)
North	-0.0439**	-0.0112	-0.000640	-0.00901	-0.0230*
	(0.0194)	(0.0183)	(0.0188)	(0.0154)	(0.0136)
Centre	-0.00766	-0.00268	-0.0122	-0.00956	-0.0107
	(0.0231)	(0.0227)	(0.0226)	(0.0185)	(0.0162)
Log cash-on-hand	-0.0113	-0.00863	-0.0255***	-0.0212***	-0.0176***
	(0.00811)	(0.00812)	(0.00826)	(0.00666)	(0.00580)
Homeowner	-0.0333*	-0.0310	-0.0199	-0.0267	-0.0270*
	(0.0197)	(0.0211)	(0.0233)	(0.0179)	(0.0150)
Implied α				0.155	0.155
				(0.041)***	(0.041)***
Observations	1,227	1,226	1,200	1,815	2,427

Notes: This table reports Tobit (marginal effect) estimates of the marginal propensity to consume (MPC), which is bounded between 0 and 1. "Only you know" is a dummy equal to one for the private, idiosyncratic bonus scenario (Scenario 2). "Public" is a dummy equal to one when the bonus is observable to coworkers (Scenarios 3 and 4), and "All receive it" is a dummy equal to one when the bonus is received by all coworkers (Scenario 4). Column 1 uses observations from Scenarios 1 and 2, with Scenario 1 as the omitted category. Column 2 uses observations from Scenarios 2 and 3, with Scenario 2 as the omitted category. Column 3 uses observations from Scenarios 3 and 4, with Scenario 3 as the omitted category. Column 4 uses observations from Scenarios 2, 3, and 4, with Scenario 2 as the omitted category. Column 5 uses observations from Scenarios 1, 2, 3, and 4, with Scenario 1 as the omitted category. Robust standard errors are reported in parentheses.

Table B3. Marginal Propensity to Transfer (MPT) heterogeneity

	<15 employees	15-100 employees	>100 employees
Public	-0.011	0.008	-0.019
	(0.027)	(0.020)	(0.019)
All receive it	-0.015	-0.006	-0.007
	(0.028)	(0.018)	(0.020)
Male	-0.003	-0.027	-0.001
	(0.023)	(0.016)*	(0.017)
Age 35 to 50	-0.005	0.000	-0.000
	(0.030)	(0.020)	(0.022)
Age 51 to 65	0.007	0.026	0.036
	(0.033)	(0.022)	(0.023)
Age 66 to 75	-0.091	-0.054	0.143
	(0.118)	(0.074)	(0.181)
Family size	0.014	0.008	0.023
	(0.010)	(0.007)	(0.007)***
College degree	0.004	-0.030	0.005
	(0.028)	(0.019)	(0.019)
North	-0.032	-0.047	-0.002
	(0.027)	(0.018)***	(0.020)
Centre	-0.036	-0.022	0.005
	(0.029)	(0.023)	(0.022)
Log cash-on-hand	-0.009	-0.004	-0.004
	(0.010)	(0.007)	(0.008)
Homeowner	0.009	-0.005	-0.052
	(0.027)	(0.020)	(0.021)**
Constant	0.168	0.133	0.118
	(0.102)	(0.069)*	(0.079)
R^2	0.03	0.06	0.04
N	216	271	500

Note. This table reports heterogeneous treatment effects of observability ("Public") and shared shocks ("All receive it") on the MPT. Columns split the sample based on firm size (number of employees). Information on number of employees in the firm is taken from a previous wave of ISCE (April 2025), and therefore the sample drops. Robust standard errors are in parentheses.

Online

Appendix C: Survey Design and Experimental Scenarios

The Italian Survey of Consumer Expectations

The Italian Survey of Consumer Expectations (ISCE) is a high-frequency household survey designed to collect detailed information on the economic circumstances, expectations, beliefs, and behavior of Italian residents aged 18–75. Conducted quarterly from October 2023 to April 2026, the survey combines a stable core questionnaire with rotating thematic modules and randomized survey experiments. It therefore provides an infrastructure both for monitoring changes in household economic conditions and for studying the causal effects of information, institutional design, and policy interventions.

The ISCE has four principal objectives: to collect frequent individual-level data on expectations and behavior; to support economic and policy analysis; to conduct survey-based experiments; and to develop and test methods for eliciting subjective expectations and beliefs. Its design draws on established international surveys, particularly the Federal Reserve Bank of New York's Survey of Consumer Expectations and the European Central Bank's Consumer Expectations Survey. It also reflects recent developments in behavioral and social-economics research that use surveys to measure perceptions, preferences, and attitudes.

Sample design and fieldwork

The ISCE sample is drawn from a proprietary panel of more than 120,000 individuals maintained by Doxa. The panel is regularly updated through both offline and online recruitment. Offline recruitment relies partly on respondents to large-scale surveys based on probabilistic face-to-face or telephone samples. Online recruitment uses multiple channels, including direct electronic communication, web advertising, and social-media campaigns. This combination is intended to broaden coverage and reduce dependence on any single recruitment method.

The survey is administered through Computer-Assisted Web Interviewing and is optimized for different devices. Approximately one-third of interviews are completed using mobile phones. Panel members receive, on average, two or three survey invitations per month. Participation is encouraged through a carefully designed incentive system in which personal rewards are accompanied by donations to nonprofit organizations. This arrangement seeks to maintain participation without attracting respondents motivated exclusively by financial compensation.

The ISCE is administered in October, January, April, and July, deliberately avoiding August and December because of their strong seasonal characteristics. A pilot survey involving 100 respondents was conducted in September 2023. It was followed by eleven regular waves, each containing approximately 5,000 completed interviews. Wave 1 was conducted in October 2023 and Wave 11 in April 2026. The resulting dataset contains more than 55,000 wave-level observations.

The sample is stratified using official population statistics provided by ISTAT. The main stratification variables are gender, age, geographical area, municipality size, education, and

employment status. Survey weights are constructed using the same variables to align the sample with the distribution of the Italian population.

Fieldwork procedures are intended to limit inattentive or excessively rapid responses. Questionnaires are randomly allocated, invitations are staggered over several days, and participation remains open for at least one week, including a weekend. These procedures reduce the overrepresentation of respondents who react immediately to invitations and allow those who check their email less frequently to participate.

The average interview lasts approximately 19 minutes. Interviews exceeding 60 minutes, around 8 percent of observations, are excluded when calculating average duration. The average response rate across the eleven waves is approximately 34 percent, with some variation over time.

The ISCE also has a substantial longitudinal dimension. Respondents who leave the panel are replaced through random replenishment, preserving the sample size while allowing repeated observation of many individuals. Of the 5,006 respondents interviewed in Wave 1, 1,694 participated in every subsequent wave through April 2026. The retention rate between consecutive waves is at least 80 percent throughout the survey. The data can therefore be used both as repeated cross-sections and as a panel for studying changes in individual expectations and behavior.

Core questionnaire

Each wave contains a stable core lasting approximately 11–12 minutes and one or more special sections lasting around four to five minutes. Experimental modules can also be incorporated by randomly assigning respondents to treatment and control groups.

The demographic and employment section records gender, age, municipality of residence, education, marital status, household size, number of income recipients, employment status, and sector of employment. Employees, self-employed workers, retirees, unemployed individuals, and job seekers are separately identified. Whenever possible, definitions and coding follow the Bank of Italy's Survey of Household Income and Wealth, facilitating comparison between the two datasets.

Income is collected in every wave.

Respondents report monthly household and individual income using eleven brackets and also indicate whether their income is well below, below, approximately equal to, above, or well above the average income of Italian households. The main variables include household disposable income, household labor and pension income, individual total income, and individual labor and pension income. All measures are net of taxes and transfers. Additional questions cover bonuses and transfers, remote working, the subjective probability of losing a job, and, for unemployed respondents, the probability of finding employment.

The wealth section measures financial wealth, real wealth, debt, homeownership, and financial-market participation. Respondents indicate whether they have transaction accounts, bonds, stocks, private pensions, life-insurance policies, health insurance, and accident insurance. Wealth and debt are reported through brackets.

The consumption section collects monthly total household expenditure, gas and electricity bills, and health expenditure. These measures allow researchers to relate consumption behavior to income, wealth, expectations, and perceived risks.

A defining feature of the ISCE is its detailed measurement of expectations. Rather than eliciting only point forecasts, respondents distribute 100 probability points across intervals representing possible future outcomes. For example, respondents report the probability that their income will fall or rise by different percentages over the following twelve months. These answers permit the construction of respondent-specific subjective probability distributions and measures of expected growth, uncertainty, and downside risk.

Expectations are collected for household disposable income, labor and pension income, total consumption, health expenditure, utility bills, house prices, returns on financial investments, GDP growth, inflation, unemployment, and mortgage interest rates. The questionnaire also measures intentions to purchase durable goods or apply for credit, perceived loan-approval probabilities, expected retirement age, expected pension replacement rates, and the likelihood of unemployment, disability, or major health expenses. Beginning with Wave 6, the order of selected expectation questions is randomized to test and limit possible order effects.

Rotating modules and experiments

The rotating modules examine a different topic in each survey wave: perceptions and expected consequences of major catastrophic risks; experiments on pension information and health insurance; use of artificial intelligence; intrahousehold decision-making; investments in energy efficiency; labor supply, and skill mismatch; opinions about poverty support, immigration, and employment protection; bargaining power; and social influences on consumption, and the design of a long-term equity fund for children.

Comparison with SHIW

Guiso and Jappelli (2026) compare sample means and medians from the ISCE with corresponding figures from the 2022 *Survey of Household Income and Wealth* (SHIW), the most recent available edition. The ISCE sample is broadly aligned with the SHIW in terms of key demographic and labor market characteristics, including gender, age distribution, employment status, and regional representation. They also present a comparison of income, consumption, and wealth statistics between the two surveys. The ISCE reports a lower median disposable income compared to the SHIW, while median consumption levels are similar across both datasets. Financial wealth appears lower in SHIW, whereas median total wealth is somewhat lower in the ISCE (136,235 vs 160,735). Financial market participation (ownership of bonds, stocks, private pensions, and life insurance) is higher among ISCE respondents.

Experimental scenarios on social interactions

Scenario 1: Unexpected Bonus – No additional information
Imagine receiving a €1,000 bonus from your employer as a one-time, unexpected payment.

Scenario 2: Unexpected Bonus – Private Information
Imagine receiving a €1,000 bonus from your employer as a one-time, unexpected payment. None of your colleagues receive this bonus: it is given only to you, and only you are aware of it.

Scenario 3: Unexpected Bonus – Public Information (Individual)
Imagine receiving a €1,000 bonus from your employer as a one-time, unexpected payment. None of your colleagues receive this bonus; it is given only to you, but everyone is aware of it.

Scenario 4: Unexpected Bonus – Public Information (Universal)
Imagine receiving a €1,000 bonus from your employer as a one-time, unexpected payment. All of your colleagues receive this bonus, and everyone is aware of it.

Q1. Allocation Question: Out of the €1,000 received as a bonus, how much would you spend, save, and give to family or friends over the next 12 months?

- Amount spent: € ___
- Amount saved: € ___
- Amount given to family/friends: € ___

Total: €1,000

Q2. Social Influence on Consumption
Imagine observing an increase in consumption (purchases, spending, vacations, etc.) among people in your social network (family, friends, colleagues, neighbors). Do you think this increase would influence your own consumption decisions?
For each of the following groups, indicate the level of influence on your consumption decisions using a scale from 1 to 10, where 1 means “no influence” and 10 means “strong influence.”

- Family: 1–10
- Friends: 1–10
- Work colleagues: 1–10
- Neighbors: 1–10