

Research Article

Testing local descriptive norms and salience of enforcement action: A field experiment to increase tax collection

Christopher Larkin*, Michael Sanders†, Isabelle Andresen*,
Felicity Algate*

Abstract: The use of behavioral science interventions, and particularly social norms, in tax compliance is a growing industry for scholars and practitioners alike in recent years. However, the causal mechanism of these interventions is unknown, where effects could be explained by a pro-social desire to support one's community, conditional cooperation, desire to conform, or fear of reprisals. We conduct a field experiment in local government taxation in the United Kingdom which tests the effectiveness of a social (descriptive) norm against a control condition and against messages that highlight the enforcement process. The social norm outperforms enforcement salience, suggesting that this explanation, although more powerful than the control, does not fully explain compliance effects. This study further provides evidence that social norm type interventions can be effective at the subnational level, a context where previous work has shown they may produce null effects.

Keywords: Taxation, Social norms, Reference groups, Subnational, Randomized controlled trial

Supplements: [Open materials](#)

Tax collection is an essential activity for government to fund its myriad public services. However, the collection process is administratively demanding. While income tax in the UK is set and collected nationally, with the administrative burden to file falling predominantly on employers, Council Tax is administered at the subnational level and taxpayers are responsible for its filing.¹ Achieving a high compliance rate (averaging 97% for Council Tax) involves a costly collection process (Ministry of Housing, Communities, and Local Government [MHCLG] 2018). If a taxpayer fails to pay by the deadline, the local authority (LA) sends them a re-

minder letter, and, where payment is not forthcoming, a second reminder letter. These reminders are followed by a court summons, hearings, and then bailiffs.² This imposes disproportionate costs on the LA and the taxpayer compared to the average size of the debt (the average annual Council Tax bill in 2018/19 for non-London areas is approximately £1,757). As collection costs increase but net revenue (i.e. the tax rate) stays the same, we see a reduction in efficiency and, in turn, total funds available for local expenditure.

In this article, we discuss a field experiment run with a medium-sized local authority in Southeast England in the United Kingdom. The experiment aimed to increase the Council Tax payment rate using two behavioral interventions delivered via reminder letters. The first letter incorporated a classic descriptive social norm (Cialdini, Reno, & Kallgren, 1990). The desire to conform to group norms is a strong motivator, and when information on majority behavior updates an individual's priors and demonstrates their behavior is in the minority, it can affect behavior change. As a large majority of debtors pay their

* The Behavioral Insights Team, office in New York, the U.S. and Manchester, the U.K.;

† What Works Centre for Children's Social Care & King's College London, the U.K.;

Address correspondence to Christopher Larkin at (chris.larkin@bi.team)

Copyright: © 2019. The authors license this article under the terms of the Creative Commons Attribution 4.0 International License.

Council Tax on time, our first intervention informed non-paying debtors that most of their peers pay promptly, and emphasized their nonconformity. Research using social norm messages (which inform taxpayers that the majority of people in their area have already paid) to increase tax compliance has been shown to work at the national level in the UK (Hallsworth, List, Metcalfe, & Vlaev, 2017) and Australia (Wenzel, 2001). Some studies using social norms to increase tax compliance have, however, failed to produce positive results (e.g., John & Blume, 2018; Homonoff, Craynor, Goldin, & Moore, 2018). This paper adds to this body of literature by using a social norm intervention at the subnational level in the UK.

For our second intervention we conceived of the consequences of non-payment as shrouded attributes of the collection process and so sought to highlight the taxpayer's proximity to enforcement. We hypothesized that taxpayers had little knowledge of the enforcement process and how close they were to being prosecuted after an initial missed payment. After failing to make a payment by the deadline, an individual is only two steps away from receiving a court summons. As this may not have been clear to a debtor, the lack of awareness could have led to their discounting the probability of legal action. We therefore designed a graphical depiction of the debt collection process and showed the debtor which stage they were currently in. For both intervention groups, we find an average increase in payment rates compared to the control, with the social norm message performing best.

This research is important for three reasons. First, it tests the use of descriptive social norms in a less commonly studied administrative context, at the subnational level, where theory predicts it will work especially well. Second, it tests descriptive norms against a more classically rational explanation of compliance behaviors. Third, it contributes to the growing use of field experiments in public administration to derive evidence for policy innovations.

The article proceeds in the following four sections. First, we describe the institutional context of subnational tax payment in the UK and explore the costs of collection for a given local authority. Second, we discuss the theoretical motivation and practical details of our interventions before, third, detailing the experimental design. Fourth, we present our results and discuss the benefits of this approach to recouping delinquent tax payments.

Institutional Context

Between 2009 and 2015, local authority budgets in England saw a decrease in their real terms spending power by an average 24.3% per person (Innes & Tetlow, 2015). This decrease has been driven by a focus on efficiency in central government and a reduction in fiscal outlays across the public sector. This in turn has led to pressure on local authorities to increase their efficiency. Despite a significant decline in spending power, local authorities remain responsible for a wide range of services including: social care, rubbish collection, road maintenance, street lighting, police, fire service, leisure centers, park maintenance, subsidizing public transport, tourism, and museums. There is considerable heterogeneity in how budgets are split across the different functions of LAs, most notably with regards to the distinct institutional arrangements of two-tier and unitary authorities in England. By far the largest part of the budget is health and social care expenditure, comprising 52% on average in 2016/17. Medway Council is a unitary authority, with competencies over council tax collection and social care, and as in many LAs, an ageing population is putting a strain on its budgets and increasing the importance of operational efficiency.

Council Tax is the primary source of locally-raised income for Local Authorities. In 2016/17, 26.0% of all budgeted revenue expenditure by LAs was raised from Council Tax (with the rest primarily made up of central government transfers). All households in an LA are liable to pay and the tax is calculated based on the 1991 valuation of a property. There are eight bands by which a dwelling can be classified, from band A, comprising the cheapest properties up to a value of \$51,254 (£40,000), to band H, comprising more expensive properties over a value of \$410,033 (£320,000). The amount a given dwelling must pay is decided by the Council in which it resides, but the national average liability for a band A property in 2017/18 was \$1,361 (£1,060) and for a band H property was \$4,086 (£3,182). Most payments are made monthly, and the burden to pay is with the household.

If a household does not make a payment following their initial bill, first and then second reminder letters are sent out, followed by a final notice and a court summons. If a household reaches the final notice stage of the process, they lose their right to pay by monthly installments and instead owe the full amount for the whole year. If an individual is summoned to court for non-payment, they must also pay

court fees and, following proceedings, possibly a \$400 (£310) enforcement agent fee. This lengthy collection process is expensive and resource intensive for the council, and some revenue from Council Tax is still not being collected. While not the focus of this paper, the introduction of non-domestic rates (NDR) – a tax on commercial properties within an LA – has recently started providing an additional and growing source of locally-raised income. NDR payments are made by businesses but, similar to the Council Tax, the burden to pay lies with the tax payer and upon missing payment a similar enforcement mechanism as outlined above is undertaken. As such the research presented here may have broader applicability beyond Council Tax collection although further research is needed specifically with NDR.

Literature Review

Previous research shows that descriptive social norms work by contrasting actual group behavior with perceived group behavior (Carrus, Bonnes, Fornara, Passafaro, & Tronu, 2009; Gerber & Rogers, 2009; Del Carpio, 2013). If people believe there is a norm of non-payment, simply informing them that a majority of people pay their tax on time highlights actual group behavior and increases the propensity to conform (i.e. pay). Some early scholarship seemed to show that the provision of social norm information had no significant effect on tax compliance behavior (Blumenthal, Christian, & Slemrod, 2001). However, more recent studies (Hallsworth et al., 2017) have shown that these kinds of interventions can be highly effective. The key difference in these studies is that the former focused on the generic case of all tax payers, who may or may not be non-compliant, whereas the latter focused on a subset of the population that were known to be non-compliant. Our study was similarly restricted to late payers.

Motivation for a social norm intervention was further justified by the local administrative context of our experiment. Prior scholarship using social norms to influence behavior has shown that closer reference groups are more effective than more distant reference groups in the presentation of norm behavior. For example, Goldstein, Cialdini, & Griskevicius (2008) showed that a closer reference group was more effective at motivating environmental conservation behaviors in hotels. Neighbors, O'Connor, Lewis, Chawla, Lee, & Fossos (2008)'s research on binge drinking behaviors among college students arrived at a similar conclusion. These findings have an

important implication for our study. Given that a local authority is a more proximal and familiar institution to residents compared to national institutions, we predict that a social norm intervention is more effective at the subnational level. This is due to: a) the degree to which an individual identifies with the other people living in their LA (as compared to the entire country), implying a more local norm; and b) a perceived greater level of scrutiny given the closeness of the institutions to the recipients. In the latter case, it is not difficult to see that being in a minority of a smaller absolute number (there were approximately 110,000 households liable to pay Council Tax in the Medway Council area in 2015, and approximately 26 million in the United Kingdom), increases the perceived level of scrutiny of individual non-compliance.

An additional motivation to use social norms at the subnational level in the UK is based on the history of Council Tax as a successor to the Community Charge tax — commonly known as the Poll Tax. The Poll Tax was introduced in England in 1990 and quickly ran into implementation problems. It was widely criticized as a regressive fiscal policy, leading to protests, riots, and campaigns to not pay (Burns, 1992). This fed a perception that it was acceptable, even desirable, not to pay the Poll Tax and that non-payment is a common behavior. The Poll Tax was replaced by Council Tax in 1993, but research has shown that some of the perception around non-payment norms persisted through to Council Tax (Dominy & Kempson, 2003; Besley, Jensen, & Persson, 2015). The extant perception connecting Council Tax to the Poll Tax, and the attitudes towards it, motivated a descriptive social norm intervention.

One other study has tested the use of a social norm intervention in a UK local authority, finding no effect on subsequent payment behavior in a first (randomized) mailing of letters and then a negative effect in a (non-randomized) rollout to all households in the LA (John & Blume, 2018). The authors note the disparity in findings between our study and theirs, highlighting three possible reasons why they might have seen a null effect and then a backfire: 1) the use of more personal language (e.g. 'you are'); 2) the fact that all households in the LA were included in their trial and not just delinquent ones (as in our study); 3) the amount of space on the letter in our study to convey and explain the social norm was greater; and 4) the difference in context between Medway and the LA where they ran their study. Medway is significantly more affluent in suburban Kent, whereas their

study was run in an inner-London borough, Lambeth, that is home to some of the highest levels of deprivation in the country and typically has a younger and more transient population. We find all of these reasons compelling. Reason (1) identifies the individual taxpayer as being a minority, thus increasing the salience of personal non-compliance. Because all households in the LA were included in their study, as described in (2), no single household owed Council Tax for that month at the point they received the letter. This compounds the likelihood that recipients believe they are not destined to be a part of a non-payer group, and certainly makes it easier to believe you will pay (even if you ultimately do not). (3) speaks to the importance of small details in how interventions are operationalized, which has been shown previously (e.g. Bryan, Walton, & Dweck, 2016). Our social norm intervention occupies central position on the letter and in relative terms comprises a large proportion of the overall text. (4) supports the broader theory on the importance of reference groups. In an area where there is a large population of highly mobile, short-term, young residents, it is possible that people feel a weaker sense of connection to their fellow residents which in turn weakens the potential impact of their reference group. In lower income communities, there is also more likely to be a higher incidence of non-payments and so there may be a stronger norm of non-payment to counteract with the intervention.

Another study that tested social norms at the subnational level showed no effect compared to a control letter (Homonoff et al., 2018). This study was run in a US state and the sample included only delinquent payers. While the sample looks more similar to our study's, in that they are already delinquent, it is possible that the much larger size of the geographic area weakened the influence of the reference group. Furthermore, as in in John & Blume (2018), the letters sent in this study included much more additional information than in ours, potentially crowding out the salience of the social norm message.

Our second intervention aimed to overcome the opacity of the enforcement process which, we believe, led people to think that non-payment would have few short or medium-term consequences. However, at the point of receiving a reminder letter, a person is only two steps away from a court summons and by highlighting this we sought to test the salience of possible enforcement action on compliance behavior. Prior scholarship has shown that perceptions of the likelihood and severity of enforcement action has an effect on tax compliance behavior (Klepper &

Nagin, 1989; Iyer, Reckers, & Sanders, 2010). Perceptions of enforcement action and tax compliance is a specific case of the broader literature around criminal deterrence (Andreoni, Erard, & Feinstein, 1998); starting with Becker (1968), an economic model of deterrence was applied to understanding criminal behavior as a classically rational phenomenon. Since this early scholarship, a large number of subsequent applications has shown that as enforcement action and perceptions of enforcement increase, the incidence of criminal behavior declines.

Allingham & Sandmo (1972) were the first to adapt this model to understanding tax compliance. Their model stipulates that tax non-compliance is motivated by the same trade-offs as any risky activity; a key implication of this being that increasing the salience of enforcement activity by locating an individual in the collection process will increase the payment rate among would-be non-compliers. This is an important counter to the fact that the consequences of non-payment are largely unobserved among peers in a community. As such, the degree of enforcement activity, its severity, or the escalating steps in this process are opaque. Homonoff et. al. (2018)'s study mentioned above also included an experimental arm that provided relatively detailed consequences of non-payment against social norms and found this to significantly increase payment rates (although at a smaller rate than our study).

The literature on social norms to date is not definitive on their precise causal pathway. Various explanations have been offered, where effects could be explained by a pro-social desire to support one's community, conditional cooperation, desire to conform, or fear of reprisals. This study builds on the existing literature by a) testing the provision of social norm information against a more classically-rational intervention that increases the salience of enforcement activity, and b) testing a social norm intervention at the subnational level, where theory predicts it will work especially well.

Methods

This experiment was a multi-arm, cluster randomized controlled trial (RCT), which ran for a three-month period in a medium-sized LA called Medway, located in the South East of England. While individual randomization is statistically efficient, cluster RCTs are often used in education and health management research settings when alternative approaches are not feasible (Puffer et al., 2005). Early exploratory work

with the LA revealed that individual randomization was not feasible, but that their Management Information System contained geographical areas called administrative units and that it was possible to randomly allocate them to receive different letter types.

There were a total of 35 units, constructed first according to the area of Medway in which they were located and then by the first letter of road names. The five areas of Medway LA were: Chatham, Gillingham, Rochester, Strood, and the parishes. This initial division and the corresponding number of administrative units for each area are shown in Figure 1 below.

The administrative units constructed for each area are then constructed based on the first letters of a household's street name. For example, administrative unit Chatham1 comprises households in Chatham whose street name starts with letters A to CA, and Chatham2 comprises households in the same geography whose road names start with letters CB to D. Because households were assigned to units using the first letter of their street address, there was little reason to think that households in the different administrative units would have different probabilities of paying their Council Tax. However, due to variation in the levels of affluence between the five areas, which has been shown to correlate with the probability of paying tax (Bloomquist, 2003), randomization was undertaken so that each administrative unit received each of the three letters (control, social norm, and enforcement salience) at some point during the trial. This meant that if an administrative unit was assigned to receive the social norm letter in month one, it would be assigned to receive the enforcement salience and control letters (in random or-

der) in months two and three respectively. The average cluster (administrative unit) size was 261 (min = 10; max = 1,717). While there is considerable variation in the size of the cluster, this did not influence the number of households assigned to each condition due to our rotating randomization strategy with each administrative unit assigned to each condition over the three months.

Importantly, each month we saw a new cohort of households within a given administrative unit become due to receive reminder letters, ensuring we never assigned the same household to receive different interventions. If a household that appeared in a previous month of the trial appeared again, they were sent the business-as-usual letter and this subsequent appearance was excluded from our analysis. This exclusion was imposed so we did not send the same household different treatment letters throughout the trial period. In addition to solving the issue of selection bias with respect to the geographies, this method of randomization had the added benefit of reducing the standard deviation through the administrative unit fixed effect by about a third, consequently increasing our statistical power.

Most households within the Medway Council area that had missed a payment were eligible to be included within the trial. There were, however, a small number of households that received a letter in error; they were eligible for an exemption but because they had just moved to the Medway LA their files had not been updated. We excluded these households from our analysis. There was also a modest amount of attrition due to some accounts closing

Figure 1
Division of Local Authority into Five Areas and 35 Administrative Units

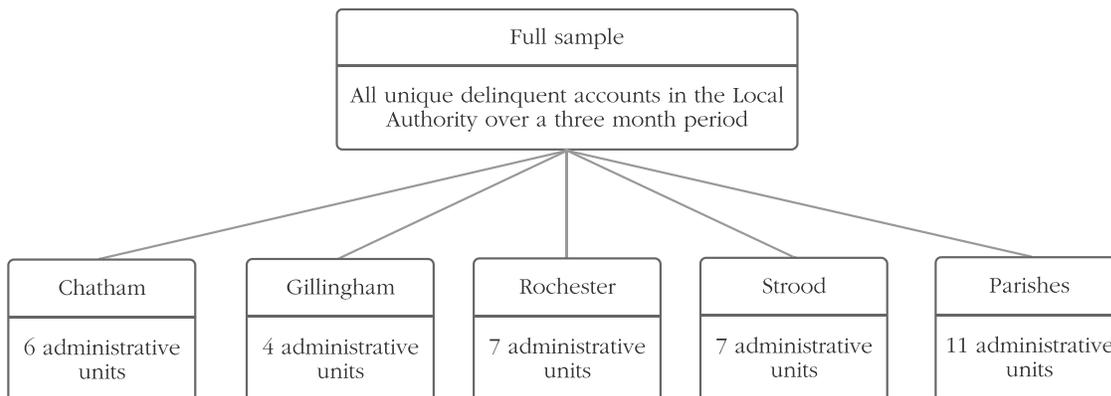


Figure 2
Assignment of Administrative Units to Trial Arm by Month

| Administrative unit | July | August | September |
|---------------------|------|--------|-----------|
| Chatham (A-CA) | 1 | 0 | 2 |
| Chatham (CE-D) | 0 | 2 | 1 |
| Chatham (E-I) | 2 | 1 | 0 |
| Chatham (J-MOS) | 1 | 0 | 2 |
| Chatham (MOU-SEC) | 0 | 2 | 1 |
| Chatham (SED-Z) | 2 | 1 | 0 |
| Gillingham (A-C) | 2 | 1 | 0 |
| Gillingham (D-K) | 1 | 0 | 2 |
| Gillingham (L-R) | 2 | 1 | 0 |
| Gillingham (S-Z) | 0 | 2 | 1 |
| Rochester (A-B) | 2 | 1 | 0 |
| Rochester (C) | 1 | 0 | 2 |
| Rochester (D-G) | 2 | 1 | 0 |
| Rochester (H-J) | 0 | 2 | 1 |
| Rochester (K-O) | 1 | 0 | 2 |
| Rochester (P-S) | 0 | 2 | 1 |
| Rochester (T-Z) | 2 | 1 | 0 |
| Strood (A-B) | 1 | 0 | 2 |
| Strood (C) | 2 | 1 | 0 |
| Strood (D-G) | 2 | 1 | 0 |
| Strood (H-L) | 0 | 2 | 1 |
| Strood (M-Q) | 2 | 1 | 0 |
| Strood (R-S) | 0 | 2 | 1 |
| Strood (T-Z) | 1 | 0 | 2 |
| All Hallows | 1 | 0 | 2 |
| Cliffe | 0 | 2 | 1 |
| Cooling | 2 | 1 | 0 |
| Cuxton | 0 | 2 | 1 |
| Frindsbury Extra | 1 | 0 | 2 |
| Halling | 0 | 2 | 1 |
| High Halstow | 1 | 0 | 2 |
| Hoo St. Werburgh | 2 | 1 | 0 |
| Isle of Grain | 1 | 0 | 2 |
| St. Mary Hoo | 2 | 1 | 0 |
| Stoke | 1 | 0 | 2 |

within the trial period. In total, there were 9,130 unique dwellings that received first reminder letters over the course of the experiment.

Depending on which administrative unit they were in and its monthly random allocation, a household was sent either the control letter, the social norm letter, or the enforcement salience letter. The control letter was the standard reminder letter Medway Council had been using (see Appendix A). The social norm letter included the phrase "*96% of Medway Council Tax is paid promptly. You are currently in the very*

small minority of people who have not paid on time" to emphasize that payment was a majority behavior (see Appendix B). And the enforcement salience letter included a flow chart that increased the salience of where in the debt collection process the recipient was located, and the subsequent legal action and costs they would be subject to if they did not pay (see Appendix C).³

We were sent panel data by Medway Council at the beginning and end of each month for the trial. The treatment identifier was constructed based on

Table 1
Treatment Allocation on Attrition and Pre-trial Household Characteristics

| TREATMENT ASSIGNMENT | (1) Attrition | (2) Benefits | (3) Arrears | (4) Tax band | (5) Disabled reduction |
|-----------------------------|--------------------------|-------------------------|------------------------|-------------------------|-----------------------------------|
| Enforcement | -0.003 (0.002) | 0.016 (0.018) | 142.51 (234.13) | -0.001 (0.070) | -0.002 (0.002) |
| Social Norm | -0.003 (0.003) | -0.006 (0.014) | 160.28 (234.30) | 0.015 (0.064) | 0.001 (0.002) |
| Constant | 0.017 (0.002) | 0.427 (0.014) | 4152.74 (166.03) | 2.365 (0.059) | 1.006 (0.001) |
| Observations | 19,127 | 9,130 | 9,130 | 9,130 | 9,130 |

Note: Standard errors clustered at the administrative unit level in parentheses

* $p < 0.5$, ** $p < 0.1$, *** $p < 0.01$

the allocation of each administrative unit in that month period. Our primary outcome variable was constructed by looking at whether an individual makes a payment between the day following the delivery of letters and the seven-day period in which a payment is required. A secondary outcome of the change in arrears was constructed by subtracting the arrears level at the end of the month from the arrears level at the beginning of the month.

Identification and Analytical Strategies

The treatment effect for a given policy intervention is defined as the average difference in individual potential outcomes (Rubin, 2005). A randomized controlled trial allows us to measure this average difference, under the condition that the random assignment allocation has been faithfully implemented and that it has generated balanced data between trial arms. The previous section speaks to the fidelity of the implementation; we now present data and the outputs from balance tests to demonstrate that randomization led to a sample balanced across a range of household-level characteristics.

Table 1 presents the results from five Ordinary Least Squares regressions testing whether treatment status significantly predicts attrition from the trial or any of a number of pre-trial household-level characteristics. These are: benefit status (a binary variable indicating whether anyone in that household received welfare benefits); arrears (in \$); tax band (eight categories from A through H); and disabled reduction (a binary variable indicating whether anyone in that household received a deduction due to a registered disability).

All p-values from the above tests are greater than 0.20 so we are confident that our sample is balanced on each of these three characteristics. Arrears level is an especially important covariate as it accounts for all debt that a given household owes the Council and so is a proxy for both length of time a payment has not been made as well as the size of payments owed. This variable captures, to some degree, the persistence of non-payment behavior for a given household. With a balanced sample across trial arms, we are confident that the models presented below provide causal estimates of our interventions' effects on payment rates and changes to arrears.

For our analysis, we pooled the three months' worth of data into one dataset and estimated a model for each of the three outcomes: 1) whether a household makes a payment for the full sample, 2) the change in a household's arrears (\$) for the subsample that made a payment, and 3) whether a household makes a payment for a restricted sample that represents a synthetic version of that seen in national social norms tax payment trials.

We used a simple OLS regression with standard errors clustered at the administrative unit level robust to heteroskedasticity.⁴ While there is a convention of using binomial logistic to analyze dichotomous outcomes, several studies have shown that fears around negative predictive values are substantiated in very few cases and that most datasets are conducive to an unbiased OLS estimate (Pohlman & Leitner, 2003; Hellevik, 2009). OLS prediction also has the added benefit of providing more intuitive estimates and standard errors. We formulate two models to look at a) the change in propensity to make a payment given treatment condition, and b) the change in payment

Table 2
Primary Treatment Effects of Letter Type on Payment Rate and Arrears

| TREATMENT ASSIGNMENT | (1) Payment rate (full sample) | (2) Δ Arrears (\$) (paying subsample) | (3) Payment rate (re- stricted sample) |
|--|--------------------------------------|--|--|
| Enforcement | 0.0688* (0.0273) | -118.27** (42.47) | 0.0784* (0.0328) |
| Social Norm | 0.1272** (0.0284) | -113.48** (43.23) | 0.1415** (0.0339) |
| Constant | 0.6297 (0.0241) | -58.37 (43.38) | 0.6140 (0.0307) |
| Fixed effects for month and administrative unit | Yes | Yes | Yes |
| Observations | 9,130 | 6,404 | 7,430 |

Note: Standard errors clustered at the administrative unit level in parentheses

* $p < 0.5$, ** $p < 0.1$, *** $p < 0.01$

amount given treatment condition. Both models take the form:

$$Y_{ia} = \alpha + \beta_1 \mathbf{T}_i + \beta_2 \mathbf{M}_{ia} + \beta_3 \mathbf{A}_i + \varepsilon_{ia}$$

*Vectors in bold

For our first model, Y_{ia} is a dichotomous variable equal to 1 if household i in administrative unit a makes a payment during the trial period, and 0 otherwise. For our second model, Y_{ia} is a continuous variable equal to the payment made by household i in administrative unit a during the trial period. α is the constant, the point where the regression line meets the y-axis; \mathbf{T}_i is a vector of binary treatment indicators, equal to 1 if household i is assigned to treatment condition T , and 0 if assigned to the control condition; \mathbf{M}_{ia} is a vector of binary indicators for the month the letters were sent out; \mathbf{A}_i is a vector of admin unit fixed effects; and ε_{ia} is an i.i.d. error term with standard errors, clustered at the administrative unit level. Where administrative unit fixed effects allow us to control for unobserved heterogeneity between units, clustering our errors at this level allows us to control for the situation where observations within each unit are not i.i.d.

Results

Table 2 shows the proportion of people who made a Council Tax payment in each trial arm within seven days of receiving their reminder letter (the deadline for a payment). The social norm and enforcement salience letters significantly increased a household's propensity to pay their Council Tax, compared to the control letter.

In the control group 62.97% of households made a payment. In the social norm group this rose to 75.69%, and for the enforcement salience group it rose to 69.84%. This was an increase of 12.72 percentage points for the social norm letter and 6.87 percentage points for the enforcement salience letter. There was also a significant difference in the payment rates between the two intervention letters ($p < .05$); that is, we can say that the social norm letter performed significantly better than the enforcement salience letter at conventional levels.

In addition to increasing the proportion of people who pay, we also saw an increase in the amount paid (conditional on paying). The treatment letters led to a significant reduction in arrears; that is, households receiving the revised letters paid more of their arrears than those receiving the control letter. A negative number indicates that the average arrears level went down throughout the trial period and is the inverse of the amount paid. Here we do not see a statistically significant difference between the social norm and enforcement salience letters, only between

each and the control group; as such we cannot make any causal claims about the differences in changes in arrears we observe across these two arms.

In relative terms, the difference between our control and social norm letters amounts to a 20.2% increase. Hallsworth et al. (2017) obtained an effect size of 1.2 percentage points (on a baseline in the control group of 8.3%) when testing a descriptive social norm message at the national level, amounting to a 14.5% change in the propensity to pay. Importantly, the effect size seen for our social norm arm is significantly higher than that seen when using a social norm intervention at the national level. This supports the theoretical construct outlined above that a descriptive social norm is more effective on Council Tax because of the unique history of the tax as well as the proximate nature of the reference group.

A cross study comparison presents more challenges to inference than a within study comparison. We address two of these here: 1) operationalization of the intervention and 2) sample differences. In terms of the intervention, the wording used in our study and Hallsworth et al. (2017)'s is almost identical.⁵ As small differences in how behavioral insights are delivered have been shown to significantly affect their impact (Bryan et al., 2016), the concordance in style and substance between the two studies is helpful in making a cross-study comparison. In terms of sample selection, it is not possible to address all bias without running a follow-up experiment that randomizes the reference group and institutional messenger (national or local government). However, one major source of potential bias lies in the average debt levels across both studies. Our analysis shows that this variable is high correlated with the likelihood of making a payment and is also a useful proxy for past payment behavior.

The average debt level in Hallsworth et al. (2017)'s sample is slightly higher than we see in this study (with a mean \approx £2,900). An implication of our theory is that if we restrict the sample to those with similar debt levels, we would see a larger effect size of the social norm intervention as compared to our entire sample. We create a crude synthetic control within our household data by isolating a comparable set of households to Hallsworth et al. (2017) that restricts the sample to have the same mean and standard deviation of debts, and by excluding debts outside of the range found in that research. To the extent that these exclusion criteria increase our relative effect size, this supports an argument that local taxes are more susceptible to social norms interventions,

suggesting a social distance argument for the effects of social norms. To the extent that the difference in the treatment effects is reduced by these exclusions, this suggests that these differences are more driven by differences between the taxpayer cohorts.

For this subsample, the effect size increases from 12.27 percentage points to 14.15 percentage points ($p < .001$). This supports a hypothesis that the effects of social norms are more powerful at a local level due to a reduction of social distance, rather than purely through updating of priors about risk.

Discussion and Conclusion

In 2014/15 financial year, Medway Council issued 139,555 initial Council Tax bills sent 31,216 first reminder letters, and collected £112,750,650 in Council Tax. Based on these figures, roughly 22.4 % of households received a first reminder letter, which equates to approximately £25,222,320 of Council Tax revenue being followed up through the successive steps of the collection process. An increase in collection rates after the first reminder letter of 12.72 percentage points, as was found by the social norm letter, would lead to an extra £2,774,320 being brought forward per year if all first reminder letters were replaced by the social norm letter. An increase in collection rates of 6.88 percentage points, as was found by the enforcement salience letter, would lead to an extra £1,841,229 being brought forward per year, if all first reminder letters were replaced by the enforcement salience letter. This represents a significant smoothing of the collection process for the local authority, where the main saving is generated through shortening the collection process and reducing the number of second reminder letters, final notices, and courts summons being distributed.

Our findings have important policy and research implications. First, they build on existing research that uses social norms to increase revenue collection and provides evidence that people are more likely to be influenced by local norms (i.e. the norm of a peer group or community) than a global norm (i.e. of a wider organization or country). By attempting to create a synthetic version of Hallsworth et al. (2017)'s sample, we partially rule out the possibility that these differential effects are driven by differences in the characteristics of tax debtors between local and national levels. Second, our findings provide indicative evidence that descriptive social norms are especially well suited to the collection of Council Tax due to the tax's unique history as a successor to the

Poll Tax. Third, our study tests the provision of social norm information against increasing the salience of enforcement and shows that the former, less classically-rational intervention is more effective at driving compliance.

We found that including the lines *"96% of Medway Council tax is paid promptly. You are currently in the very small minority of people who have not paid on time"* in a notice of non-payment letter was effective at increasing tax compliance among late payers of Council Tax. This intervention provided social norm information (e.g. that most people do in fact pay on time) to non-payers and generated a 20% increase in the payment rate. Importantly, we tested this intervention against a different, more classically-rational letter that highlighted where an individual debtor was located in the collection process, emphasizing they were only two steps away from a court hearing. While the intervention highlighting the enforcement consequences of continued non-payment was more effective than the control letter, the effect was significantly smaller than in the social norm arm. Future research could look at different ways of how enforcement salience is presented as a way to motivate tax compliance. In our study we used a flow diagram that presented enforcement as four steps, using equal-sized boxes for each (see Appendix B for a copy). An important omission from this operationalization is that we provide no information to recipients about the probability of progressing to each step.

Another important contribution of this research is that it shows that social norms can be more effective at the subnational level, and perhaps particularly so in the UK due to the unique history of the Council Tax. By constructing a synthetic control within our sample of Medway residents and comparing the differential effect of the social norm treatment letter in this sample with the effects seen in other scholarship using similar interventions but at the national level, we showed that the treatment effect is larger when the norm comes from a more local reference group.

Given the null effects we've recently seen for social norm interventions (and, importantly, one at the subnational level in the UK), future research should do more to identify when and why we may see positive effects. One possible line of research could explore the importance of norm salience with respect to other information presented to residents. Another line of research could focus on whether the behavior that is being identified as minority behavior is prospective or established. In this study, and others

where social norms have been effective, the intervention was sent only to delinquent households (i.e. the minority behavior was a fact on the ground), whereas studies that have seen null effects have included all households. While this study was not designed to test the interaction between prospective versus established minority behaviors and the salience of norm exposition within an intervention, we believe this remains an important question in the field.

Acknowledgment

We are grateful to Medway Council for their support running this trial, especially Patrick Knight. We are also grateful to the two anonymous reviewers, and editors at the Journal of Behavioral Public Administration, whose comments and edits on earlier versions of this manuscript have made it immeasurably better. As ever, all mistakes remain with the authors.

Notes

1. Pay As You Earn (PAYE) was introduced in 1944 due to war time pressures on British public finances. As a more efficient system was needed to collect more tax from more people, Sir Paul Chambers and Sir Cornelius Gregg devised the pay-as-you-earn system where businesses became liable to pay tax on behalf of their employees. In the UK the majority of (national) income tax is paid via PAYE, meaning that for this tax most people need not take any action themselves. In 2014, approximately two thirds of all taxpayers pay via the PAYE system, accounting for approximately 90% of all income tax (Pope & Roantree 2009).
2. Bailiffs in the UK are usually private contractors that execute the decisions of a court. In this context, they would go to debtors' properties to retrieve monies owed or property that equals the same value.
3. There are other cosmetic differences between the treatment and control letters. However, these do not relate to the novel contribution of this article which lies in testing social norms against enforcement salience and in comparing the effect of social norm interventions in a more local sample against other national samples where the same type of intervention has been studied.

4. We also specified the same models clustering the errors using a (month*administrative unit) variable. All results reported below are robust to this specification.
5. Hallsworth et al. (2017)'s study used the phrase: "Nine out of ten people in the UK pay their tax on time. You are currently in the very small minority of people who have not paid us yet."

References

- Allingham, M., & Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3–4): 323–338.
- Andreoni, J., Erard, B., & Feinstein, J. (1998). Tax compliance. *Journal of Economic Literature*, 36(2): 818–860.
- Becker, G. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2): 169–217.
- Besley, T., Jensen, A., & Persson, T. (2015). Norms, enforcement, and tax evasion. 10372, Centre for Economic Policy Research.
- Bloomquist, K. (Ed.) (2003). Tax evasion, income inequality, and opportunity costs of compliance., vol. 96. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association.
- Blumenthal, M., Christian, C., & Slemrod, J. (2001). Do normative appeals affect tax compliance? Evidence from controlled experiment in Minnesota. *National Tax Journal*, 54: 125–138.
- Bryan, C., Walton, G., & Dweck, C. (2016). Psychologically authentic versus inauthentic replication attempts. *Proceedings of the National Academy of Sciences*, 113(26): 7112–7117.
- Burns, D. (1992). *Poll Tax Rebellion*. London: AK Press.
- Carrus, G., Bonnes, M., Fornara, F., Passafaro, P., & Tronu, G. (2009). Planned behavior and "local" norms: An analysis of the space-based aspects of normative ecological behavior. *Cognitive Processing*, 10: 198–200.
- Cialdini, R., Reno, R., & Kallgren, C. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6): 1015–1026.
- Del Carpio, L. (2013). *Are the neighbors cheating? Evidence from a social norm experiment on property taxes in Peru.*, Princeton University.
- Dominy, N., & Kempson, E. (2003). Can't pay or won't pay: A review of creditor and debtor approaches to the non-payment of bills. 4/03, Personal Finance Research Centre, University of Bristol.
- Gerber, A., & Rogers, T. (2009). Descriptive social norms and motivation to vote: Everybody's voting and so should you. *The Journal of Politics*, 71: 178–191.
- Goldstein, N., Cialdini, R., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35: 472–482.
- Hallsworth, M., List, J., Metcalfe, R., & Vlaev, I. (2017). The behavioralist as tax collector: using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148: 14–31.
- Hellevik, O. (2009). Linear versus logistic regression when the dependent variable is a dichotomy. *Quality and Quantity*, 43: 59–74.
- Homonoff, T., Craynor, T., Goldin, J., & Moore, L. (2018). Communicating tax penalties to delinquent payers: Evidence from a field experiment. *Unpublished Manuscript*. Available here: <https://drive.google.com/file/d/1CBT-YilYChexF3Yk21WH8pGcJZexL1VS/view>
- Innes, D., & Tetlow, G. (2015). Central cuts, local decision-making: Changes in local government spending and revenues in England, 2009-10 to 2014-15. Briefing Note BN166, Institute for Fiscal Studies.
- Iyer, G., Reckers, P., & Sanders, D. (2010). Increasing tax compliance in Washington state: A field experiment. *National Tax Journal*, 63: 7–32.
- John, P., & Blume, T. (2018). How best to nudge taxpayers? The impact of message simplification and descriptive social norms on payment rates in a central London local authority. *Journal of Behavioral Public Administration*, 1(1): 1-11.
- Klepper, S., & Nagin, D. (1989). Tax compliance and perceptions of the risks of detection and criminal prosecution. *Law & Society Review*, 23: 209–240.
- Ministry of Housing, Communities, and Local Government. (2018). Collection rates and receipts of council tax and non-domestic rates in England 2017-18 – revised. *Statistical release*.
- Neighbors, C., O'Connor, R., Lewis, M., Chawla, N., Lee, C., & Fossos, N. (2008). The relative impact of injunctive norms on college student drinking: The role of reference group. *Psychology of Addictive Behaviors*, 22: 576–581.
- Pohlman, J., & Leitner, D. (2003). A comparison of ordinary least squares and logistic regression. *The Ohio Journal of Science*, 103: 118–125.
- Pope, T., & Roantree, B. (2009). A survey of the UK tax system. Briefing Note BN09, Institute for Fiscal Studies.
- Puffer, S., Torgerson, D., & Watson, J. (2005). Cluster randomized controlled trials. *Journal of Evaluation in Clinical Practice*, 11: 479–483.
- Rubin, D. (2005). Causal inference using potential outcomes: Design, modelling, decisions. *Journal of the American Statistical Association*, 100: 322–331.
- Wenzel, M. (2001). Misperceptions of social norms about tax compliance (2): A field experiment. Working Paper No.8, Centre for Tax System Integrity.