

A Field Experiment: Testing the Potential of Norms for Achieving Behavior Change
in English Parishes

Julie Van de Vyver^{a*} & Peter John^b

^a School of Psychology, University of Lincoln, UK

^b School of Public Policy, University College London, UK

This is the submitted version of the article which has been accepted for publication in the *Journal of Applied Social Psychology* (2017). For a copy of the final manuscript see [http://onlinelibrary.wiley.com/journal/10.1111/\(ISSN\)1559-1816](http://onlinelibrary.wiley.com/journal/10.1111/(ISSN)1559-1816)

*Correspondence concerning this article should be addressed to Dr Julie Van de Vyver, School of Psychology, University of Lincoln, Lincoln, UK. E-mail: jvandevyver@lincoln.ac.uk

Abstract

Tests of behavioral insights using field experiments have become more common in recent years, and have been deployed by UK government and its agencies. Typically these field experiments aim to change individual-level behaviors. The current paper tests the potential of behavioral insights for changing group-level behavior. This paper reports the results of a field experiment carried out with the Department of Communities and Local Government. The field experiment tested whether a normative message (vs. a neutral or no message) could encourage parish councils to register an asset of community value (social action). There was no statistically significant effect from this intervention, but the process of designing and implementing this field experiment shows the potential for theories of behavior change to be used by government departments.

A Field Experiment: Testing the Potential of Norms for Achieving Behavior Change
in English Parishes

Behavioral Insights Revolution

In recent years there has been considerable official interest in the use of *behavioral insights* to inform public policy. Policy-makers are starting to take advantage of the wealth of behavioral research from across areas of psychology, economics, and other social sciences, which has been popularized in recent years (see Thaler & Sunstein, 2008). Behavioral insights (or nudge) is a term that is increasingly used and understood across disciplines as well as by lay-people. Behavioral insights typically consist of low-cost strategies that aim to change behavior. The years since 2010 have seen central government departments and agencies in the UK and elsewhere using behavioral insights to complement conventional policy tools (i.e., legislation, regulation, incentives) and to enhance policy outcomes (see Dolan et al., 2012). Examples of behavioral insights interventions that were first tested using experimental designs in the field and then rolled out to improve public policy include using tax reminders to collect taxes and using SMS texts to collect court fines (Halpern, 2015).

Governments have always sought to influence human behavior and such an interest goes back to recruiting for armies and getting citizens to pay their taxes, core functions of government. The study of human behavior influenced the emergence of central regulation in the nineteenth century, such as the rise of public health legislation and regulation, and continued in the twentieth as government became more professionalized and responsive to specialist knowledge, such as measures to prevent traffic accidents. However, the greater focus on behavior change as a key objective of government emerged more strongly in the twenty-first century following on from the

academic advances in the social sciences. Central government in the UK has taken a great interest in this agenda and it has influenced central government discussion documents (e.g. Halpern, Bates, Mulgan, & Aldridge, 2004). But the visibility of behavioral sciences increased markedly with the publication of *Nudge* (Thaler & Sunstein, 2008). Most prominent has been the emergence of the Behavioral Insights Team (BIT), established in 2010 under the United Kingdom Coalition government, which succeeded in gaining public attention and a measure of success, and was spun out of government as an independent organization (Halpern, 2015). Key has been the use of the randomized controlled trials (or field experiments) as a way to test the effectiveness of these behavioral interventions in the field (Haynes, Service, Goldacre, & Torgerson, 2012; John, 2013). These have been done on standard operating procedures of government and on new policies, which are intended to show the potential for a greater roll out.

Influencing Organization-Level Behavior

The current research will examine the potential of behavioral insights for influencing organization-level behavior among parish councils. This is a particularly important research endeavor for several reasons. First, while ample research (both in the laboratory and in the field) has examined the effects of behavioral insights on individual-level behavior, much less is known about whether and how behavioral insights may affect group-level behavior. Indeed, in spite of the diffusion of behavioral insights practices and the greater use of scientific knowledge, English local and parish councils have experienced fewer of these behavioral insights interventions. It seems likely that organization-level behavior may be harder, or at least more complex, to influence than individual-level citizen behavior. Indeed, action from the organization is dependent on the decision-making structure within the organization. In

other words, organization-level action may be subject to the decision of veto players in the organization or may suffer from implementation loss through chains of command or simple failure of collective action. Nevertheless, organizations are headed by individuals who can respond to behavioral cues. There are examples of the use of behavioral interventions directed to organizations in the less developed context, such as encouraging the standing of women in village councils in India (Chattopadhyay & Duflo, 2004) or linking aid to performance with local providers in India (Olken, Onishi, & Wong, 2014). In the UK, behavioral interventions targeted at businesses, such as Growth Vouchers, show how policy-makers can seek to influence organizations (Department for Business and Skills and Cabinet Office 2014).

Second, given that UK governments elected since 2010 have given a particular emphasis to the philosophy of localism, which implies a hand-off approach to local government (Lowndes & Pratchett, 2012), novel strategies need to be developed and tested to help advance knowledge and ensure better implementation of policies among local government organizations. Effective behavioral insights interventions might be seen as consistent with the localism agenda as such interventions do not use the authority of the state to command local authorities but instead indirectly encourage behaviors.

In sum, the current research aims to contribute to the literature by testing the effectiveness of a behavioral insights intervention for influencing group-level (rather than individual-level) behavior in a relevant applied context (i.e., among parish councils).

Use of Descriptive Norms

In this experiment, we draw on the behavioral insights concept of *social norms* (see Fehr & Fischbacher, 2004), which has become the tool of choice of today's

behavioral policy-makers. Social norms can be understood as socially shared definitions of the way people do behave or should behave (Fehr & Fischbacher, 2004). Research on normative theory has suggested that two types of social norms exist – descriptive and injunctive (Cialdini, Kallgren, & Reno, 1991; Fehr and Fischbacher, 2004). Descriptive norms are those that “characterize the perception of what most people do” (Cialdini et al., 1991, p. 203). In contrast, injunctive norms are those that “characterize the perception of what most people approve or disapprove [of]” (Cialdini et al., 1991; p. 203). For example, the statement “nine out of ten people pay their tax on time” would be a descriptive norm. In contrast, the statement “nine out of ten people think that people should pay their tax on time” would be an injunctive norm.

People are heavily influenced both by what other people do and by what other people think they should do. Social norms offer a standard from which people do not want to deviate as they do not want to risk social exclusion from their group (Schultz et al., 2007). Therefore, an effective strategy for mobilizing people to act is to publicize the desired social norm. This strategy is a very well validated in the psychology and economics literatures and has become a popular tool among policy-makers to encourage compliance (Dolan et al., 2012).

Empirical evidence suggests that social norms can mobilize people to act. Indeed, research has drawn on social norms to change socially important behaviors such as alcohol consumption, drug use, disordered eating, gambling, littering, and recycling (e.g., Donaldson, Graham, & Hansen, 1994; Larimer & Neighbors, 2003; Neighbors, Larimer, & Lewis, 2004; Schultz, 1999; Schultz, Tabanico, & Rendon, 2008; see also Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007). For example, one experiment tested whether different types of messages could improve

the reuse of towels in hotel rooms. A sign that asked guests to reuse towels in order to save the environment led to 37% of people doing so. A sign that highlighted that most other guests in the hotel reused their towel more than once led to 44% of people reusing their towels. A sign that highlighted that most occupants of the same room reused their towels more than once led to 49% of people reusing their towels (Goldstein, Cialdini, & Griskevicius, 2008). Therefore highlighting a social norm significantly increased pro-environmental behavior in a field experiment.

The Current Research

The current research focuses on the potential of descriptive norms for encouraging parish councils to act. Specifically, in the current research we explore whether a descriptive norm message (vs. a neutral message or no message) can increase parish councils' uptake of community rights policies, in particular the *Community Right to Bid*. The community right to bid gives parish councils and community groups the power to protect any local building or land by nominating it as "an asset of community value". The advantages of nominating a community asset, if the nomination is accepted by the local authority, are that (1) the parish council (or community group) will be notified if the asset is ever put up for sale; (2) the parish council (or community group) will have the opportunity to buy it if they wish to do so; (3) a signal is sent to the owner making them aware that the asset is valued by the community; and (4) the nomination could be taken into consideration by planning authorities when considering planning applications (see Department for Communities and Local Government, 2015). During the development stage of this study, 1002 community assets (e.g., community halls, pubs) were listed across the UK.

In the current experiment parish councils were sent either an e-mail depicting the descriptive norm, an identical e-mail but without the descriptive norm, or no e-

mail at all. Parish council's uptake of community rights (i.e., asset nomination) was subsequently monitored. Based on past research (e.g., Schultz et al., 2007), it is hypothesized that asset nomination will be higher among those who viewed the descriptive norms message than among those who viewed the neutral message or among those who received no message at all. Furthermore, it is hypothesized that asset nomination will be higher among those who viewed the neutral message than among those who received no message at all. This is because, while unlike the experimental norm-Condition there is no motivational incentive, the neutral message provides parishes with easy, accessible, and necessary information about the community right. Indeed, research suggests that while motivational incentive is important for behavioral change, so is capability and opportunity (i.e., knowledge and resources) (see Michie, van Stralen, & West, 2011).

Method

Sample

The sample consisted of 3,026 parish councils from all regions in England except for London (see Online Supplemental Materials for a full breakdown of the sample by regions, counties, and local authorities). The sampling strategy involved scanning the web for open-access contact details (e-mail addresses) of parish councils - typically accessed via Local Authority webpages. The sample was therefore restricted by public availability of parish council contact information. The mean precept size (i.e., parish income from taxes paid by residents within the parish area) for this sample was £27, 599.61 ($SE = 924.22$). Precept size ranged from £48 to £801, 462. Twenty-nine per cent of parishes in this sample had a precept of less than £5,000 (vs. 34 % of parishes across England). Ninety-three per cent of parishes in this sample

had a precept of less than £100,000 (like 90 % of parishes across England; Department for Communities and Local Government, 2014).

Design and Stimuli

Parishes were randomly assigned to Condition in a one factor between participants design with three levels (Message: norms, neutral, none).¹ The norm and neutral messages were distributed via e-mail (using Mailchimp, a web-based e-mail marketing application). The e-mails were sent by the Department for Communities and Local Government. Messages (including message headers) were kept identical across Conditions. Titles within the e-mail varied by Condition (see Table 1 for a depiction, see also Online Supplemental Materials).

Randomization. In order to randomly allocate parish councils to one of the three Conditions, block randomization was employed. Block randomization (randomizing within subgroups) ensured (1) a relatively equal number of small versus large parishes (in terms of precept size) and (2) a relatively equal number of parishes with (vs. without) assets already listed across the three Conditions.

A phenomenon that became apparent while retrieving parish councils' contact details was that parish councils sometimes shared the same clerk (i.e., administrator) with other parish councils (25.4% of the parish councils in the dataset shared a clerk with at least one other parish council). Therefore, in the current experiment, one parish council was randomly selected per clerk. However, while only one parish per clerk was randomly assigned to Condition, the behavior of all parishes will be analyzed. It is highly likely that the clerk would share the information with all of the parish councilors who they work with. This inclusion of the duplicate clerks in the analyses increases the sample size and therefore the statistical power.²

Table 1

Message Headers and Titles

Condition	E-mail header	E-mail title
1.No e-mail	N/A	N/A
2.Neutral e-mail	Protect a community asset – it's quick and easy	Protect a community asset – it's quick and easy
3. Norm e-mail	Protect a community asset – it's quick and easy	Parishes and local groups all over the country are nominating community assets. Over 1000 assets have already been listed.

Measures

Outcomes were measured in two ways: (1) monitoring e-mail activity; and (2) monitoring community asset listings.

Interest in listing an asset. Mailchimp software was used to send out the neutral and norm messages. Mailchimp is able to monitor whether parish councils clicked on the link for further information given at the end of the e-mail (see Online Supplemental Materials). This offers some insight into whether the norm message generated greater, lesser, or equal interest in listing a community asset (vs. the neutral message). No reminder e-mail was sent.

Asset listings (behavioral outcome). We monitored whether parish councils listed community assets (offering behavioral data). Listings were monitored via inspecting local authority webpages. Listings made prior to the intervention (i.e., prior to April 2014) will be referred to as Time 1 listings. Listings made between April and August 2014 will be referred to as Time 2 listings. Due to the nature of the collaboration with DCLG, only data up to August 2014 were available for analyses.

Nevertheless, the four-month period seemed a reasonable length of time to allow parish councils to nominate assets of community value.

Covariates. We also collated information about: (1) geographical area of the parish council, (2) whether a community asset was already listed in the boundaries of the parish council, and (3) precept size of the parish. These three variables will be used as covariates in the analyses on the grounds that areas vary considerably across England, that prior listing might affect responsiveness, and that the financial capacity would affect the capability of responding.

Results

For the norm and neutral Conditions, data analyses were restricted to those who opened the e-mail, and therefore to only those who were exposed to Condition. In the norms condition 437 (43.8%) parish councils opened the e-mail. In the neutral condition 433 (43.4%) parish councils opened the e-mail.

Interest in Listing an Asset

Statistical analyses of the Mailchimp data were conducted three weeks after the e-mails were sent. Binary logistic regression was conducted to examine whether Condition (norms message vs. neutral message) predicted whether or not parish councils clicked on the link for further information. Region of parish, asset listed, and precept size were also entered into the model as predictors.

Results showed that Condition did not significantly predict whether parishes clicked on the link for further information ($B = -.22$, $SE = 0.26$, $p = .397$, $\text{ExpB} = .80$). Precept size was the only significant predictor ($B = .32$, $SE = 0.11$, $p = .002$, $\text{ExpB} = 1.38$). Specifically, the greater the precept size, the more likely parishes were to click on the link for further information.

Asset Listings (Behavioral Outcome)

Whether a new asset was listed. An analysis of variance was conducted to test the effect of Condition (norm message vs. neutral message vs. no message) on whether a new asset was listed (0 = *no*, 1 = *yes*). Analyses controlled for region of parish and precept size. Results revealed no significant main effect of Condition, $F(2, 2413) = 0.65, p = .522, \eta^2 = .001$ (see Table 2). The only significant effect was a main effect of precept size, $F(1, 2413) = 4.56, p = .033, \eta^2 = .002$.

Number of assets listed. A mixed model analysis of variance was conducted to test the effect of Condition on the number of assets listed at Time 1 versus at Time 2. Analyses controlled for region of parish and precept size. We employed a 2 (number of assets listed at Time 1 vs. number of assets listed at Time 2) x 3 (Condition: norm message vs. neutral message vs. no message) design with the first factor being a within participants factor and the second factor being a between participants factor. Results revealed no significant main effect of assets listed, $F(1, 2413) = 0.06, p = .805, \eta^2 < .001$, no significant main effect of Condition, $F(2, 2413) = 0.92, p = .400, \eta^2 = .001$, and no significant two-way interaction, $F(2, 2413) = 0.48, p = .620, \eta^2 < .001$ (see Table 2). The only significant interaction, echoing the results above, was a Precept size x asset listed interaction, $F(1, 2413) = 9.86, p = .002, \eta^2 = .004$.

Table 2

Means and Standard Errors for the Effect of Condition on Assets Listed

Condition	Whether New	Assets	Assets
	Asset Listed	Listed at Time 1	Listed at Time 2
	<i>M (SE)</i>	<i>M (SE)</i>	<i>M (SE)</i>
Norms message	.04 (.01)	.07 (.02)	.13 (.03)
Neutral message	.04 (.01)	.05 (.02)	.09 (.03)
No message	.03 (.01)	.07 (.01)	.13 (.02)

Note. “Whether new asset listed” ranged from 0 = *no new asset listed* to 1 = *new asset listed*, therefore the means can be interpreted as proportions. “Assets listed at Time 1” reflected the number of assets listed at Time 1 and ranged from 0 to 9. “Assets listed at Time 2” reflected the number of assets listed at Time 2 and ranged from 0 to 15.

Discussion

This field experiment was designed to test the effectiveness of a behavioral insights technique (i.e., norms) to encourage parish councils to use the Community Right to Bid (i.e., to nominate assets of community value). The results showed that there was no significant effect of Condition (norms message vs. neutral message vs. no message) on assets listed at a four-month follow-up.

Overall the experiment showed how the behavioral sciences can be adapted to seek behavior change in the decentralized context with a valid methodology. The nudge approach was light-touch and consistent with a hands-off approach to implementing policies in the locality. The results highlight the importance of experimentally testing the effectiveness of behavioral insights techniques before

further roll out. Indeed, the effectiveness of behavioral insights techniques will depend on context, audience, methodology, and many other factors. Therefore, it is essential for policy makers to work with researchers to conduct controlled field experiments prior to investing substantial time and money in such initiatives.

The non-significance of the findings could have occurred for many reasons. First, it is possible that the descriptive norm intervention simply was not effective at motivating behavioral change. However, given the substantial amount of prior evidence to the contrary (e.g., Donaldson, Graham, & Hansen, 1994; Larimer & Neighbors, 2003; Neighbors, Larimer, & Lewis, 2004; Schultz, 1999; Schultz, Tabanico, & Rendon, 2008; see also Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007) this explanation seems unlikely, or at least insufficient. Second, it is possible that organization-level behavior is more difficult to influence than individual-level behavior due to decision-making structures within the organization or implementation loss through chains of command. Third, it is possible that local parishes are resistant to encouragements from the centre. Future research should explore these possible reasons further using quantitative and/or qualitative research to explore parish councilors' attitudes, motivations, and intentions on this topic.

This field experiment had some limitations, which is inherent to any intervention of this nature. One such issue was the availability of the asset nomination data. Specifically, parish councils' asset nominations needed approval by the local council in order to be successfully listed. Given that we were interested in parish councils' behavioral change, we were interested in their *nominations*. Nevertheless, the dependent variable in the analyses reflected successful listings rather than nominations. This was due to data availability restrictions. Analyzing the successful listings (rather than nominations) means that (1) we cannot be certain that a

nomination was not made before the intervention occurred, and (2) we cannot be certain that unsuccessful (or delayed) nominations were not made following the intervention. Nevertheless, the random assignment to groups helps attenuate these issues. The key suggestion for future research in this particular field is to conduct a similar field experiment but over a longer period of time. This would ensure more accuracy in the asset listing data. It would also be interesting to explore the effects of different behavioral insights techniques. Due to the sample size we were only able to test one behavioral insights technique, but it is very plausible that other behavioral insight techniques could be more effective as compared to the descriptive norm technique employed in the current experiment. Finally, it would be interesting to explore whether longer-term or repeated exposure to the behavioral insights message(s) would be more effective than a one-off exposure. As local parish and town councils become more important and better known by central government, it may be possible to craft other light-touch interventions that communicate government policy in a way that is suitable to the local context and appropriate for local democratic choice.

References

- Chattopadhyay, R., & Duflo, E. (2004). Women as policy makers: Evidence from a randomized policy experiment in India. *Econometrica*, *72*, 1409-1443.
doi:10.1111/j.1468-0262.2004.00539.x
- Cialdini, R. B., Kallgren, C. A., & Reno, R. R. (1991). A focus theory of normative conduct: A theoretical refinement and reevaluation of the role of norms in human behavior. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol 24, pp. 201–234). New York, NY: Academic Press.
- Department for Business and Skills and Cabinet Office. (2014). *Growth vouchers trial protocol*. London: Department for Business and Skills and Cabinet Office.
Retrieved via: <http://www.behaviouralinsights.co.uk/wp-content/uploads/2015/07/bis-14-561-growth-vouchers-programme-trial-protocol.pdf>.
- Department for Communities and Local Government (2014). *Parishes and other local precepting authorities: 2014-15 England*. London: Department for Communities and Local Government. Retrieved via:
https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/335854/Council_Tax_Levels_set_by_local_precepting_authorities_in_England_2014-15__3_.pdf
- Department for Communities and Local Government (2015). *Policy paper: 2010 to 2015 government policy: Localism*. London: Department for Communities and Local Government. Retrieved via:
<https://www.gov.uk/government/publications/2010-to-2015-government-policy-localism/2010-to-2015-government-policy-localism>

- Dolan, P., Hallsworth, M., Halpern, D., King, D., Metcalfe, R., & Vlaev, I. (2012). Influencing behaviour: The mindspace way. *Journal of Economic Psychology, 33*, 264-277. doi:10.1016/j.joep.2011.10.009
- Donaldson, S. I., Graham, J. W., & Hansen, W. B. (1994). Testing the generalizability of intervening mechanism theories: Understanding the effects of adolescent drug use prevention interventions. *Journal of Behavioral Medicine, 17*, 195-216. doi:10.1007/BF01858105
- Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior, 25*, 63-87. doi:10.1016/S1090-5138(04)00005-4
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research, 35*, 472-482. doi:10.1086/586910
- Halpern, D. (2015). *Inside the Nudge Unit: How small changes can make a big difference*. London: Penguin Random House.
- Halpern, D., Bates, C., Mulgan, G., & Aldridge, S. (2004). *Personal responsibility and changing behaviour: The state of its knowledge and its implications for public policy*. London: Cabinet Office, Prime Minister's Strategy Unit.
- Retrieved via:
<http://webarchive.nationalarchives.gov.uk/+http://www.cabinetoffice.gov.uk/media/cabinetoffice/strategy/assets/pr2.pdf>
- Haynes, L., Service, O., Goldacre, B., & Torgerson, D. (2012). *Test, learn, adapt: Developing public policy with randomised controlled trials*. London: Cabinet Office. Retrieved via:

https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/62529/TLA-1906126.pdf

John, P. (2013). Experimentation, behaviour change and public policy. *The Political Quarterly*, 84, 238-246. doi:10.1111/j.1467-923X.2013.12010.x

Larimer, M. E., & Neighbors, C. (2003). Normative misperception and the impact of descriptive and injunctive norms on college student gambling. *Psychology of Addictive Behaviors*, 17, 235-243. doi:10.1037/0893-164X.17.3.235

Lowndes, V., & Pratchett, L. (2012). Local governance under the coalition government: Austerity, localism and the 'Big Society'. *Local Government Studies*, 38, 21-40. doi:10.1080/03003930.2011.642949

Michie, S., van Stralen, M. M., & West, R. (2011). The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation Science*, 6, 42. doi:10.1186/1748-5908

Neighbors, C., Larimer, M. E., & Lewis, M. A. (2004). Targeting misperceptions of descriptive drinking norms: efficacy of a computer-delivered personalized normative feedback intervention. *Journal of Consulting and Clinical Psychology*, 72, 434-447. doi:10.1037/0022-006X.72.3.434

Olken, B. A., Onishi, J., & Wong, S. (2014). Should aid reward performance? Evidence from a field experiment on health and education in Indonesia. *American Economic Journal: Applied Economics*, 6, 1-34. doi:10.1257/app.6.4.1

Schultz, P. W. (1999). Changing behavior with normative feedback interventions: A field experiment on curbside recycling. *Basic and Applied Social Psychology*, 21, 25-36. doi:10.1207/s15324834basp2101_3

Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V.

(2007). The constructive, destructive and reconstructive power of social

norms. *Psychological Science*, *18*, 429–434. doi:10.1111/j.1467-

9280.2007.01917.x

Schultz, P. W., Tabanico, J., & Rendon, T. (2008). Normative beliefs as agents of

influence: Basic process and real-world applications. In R. Prislin & W. Crano

(Eds.), *Attitudes and persuasion* (pp. 385–409). New York: Psychology Press.

Thaler, R. H., & Sustein, C. R. (2008). *Nudge: Improving decisions about health,*

wealth and happiness. New Haven: Yale University Press.

Footnotes

¹ E-mails were sent out on Wednesday 9th and Thursday 10th April 2014. Due to a procedural error using Mailchimp the neutral messages were sent out 24 hours before the norm messages. This was not intended and we appreciate that this means allocation of Condition is no longer entirely random but influenced by time. However, we think it unlikely that one day's difference would make a difference to the impact of the treatment.

² The results do not differ significantly by whether the duplicate clerks were included or excluded.

Acknowledgments

This research was supported by the Economic and Social Research Council (ES/J500148/1). We thank the Department for Communities and Local Government for facilitating this research. We particularly appreciate the contribution of Dr Maria O'Beirne and Tom Tolfree. The contents of this paper do not reflect the policy or views of the Department for Communities and Local Government or those of the Government Social Research Service.