**British Journal of
Political Science**

LETTER

Unequal and Unsupportive: Exposure to Poor People Weakens Support for Redistribution among the Rich

Matias Engdal Christensen[1] (iD), Peter Thisted Dinesen[2,3] (iD) and Kim Mannemar Sønderskov[1] (iD)

[1]Department of Political Science, Aarhus University, Aarhus, Denmark, [2]Department of Political Science, University of Copenhagen, Copenhagen, Denmark and [3]Department of Political Science, University College London, London, UK
Corresponding author: Matias Engdal Christensen; Email: mec@ps.au.dk

Abstract

Do the rich become more or less supportive of redistribution when exposed to poor people in their local surroundings? Most existing observational studies find that exposure to poor individuals is positively associated with support for redistribution among the well-off, but one prominent field experiment found a negative link. We seek to resolve these divergent findings by employing a design closer to the studies that have found a positive link, but with more causal leverage than these; specifically, a three-wave panel survey linked with fine-grained registry data on local income composition in Denmark. In within-individual models, increased exposure to poor individuals is associated with lower support for redistribution among wealthy individuals. By contrast, between-individual models yield a positive relationship, thus indicating that self-selection based on stable individual characteristics likely explains the predominant finding in previous work.

Motivated by concerns over growing income inequality across developed democracies (Piketty and Saez 2014), a large and growing literature has investigated the relationship between inequality and voters' support for redistribution (Condon and Wichowsky 2020; Kenworthy and McCall 2008; Lupu and Pontusson 2011). The nature of this relationship is crucial because voters can—via their policy preferences that animate politicians to enact policies—contribute to both the amelioration (Meltzer and Richard 1981) and the aggravation of economic inequality (Kelly and Enns 2010) depending on how they respond to it.

A significant body of work suggests that the link between income inequality and support for redistribution is partly rooted in local experiences. Economic differences observed in one's local surroundings shape perceptions of economic inequality and ultimately influence attitudes toward redistribution (Condon and Wichowsky 2020; Newman 2020; Sands 2017; Sands and De Kadt 2020).[1] However, there is disagreement over the empirical evidence.

---

[1]In the literature, the term 'local economic inequality' is used to refer to both aggregate measures of the local income distribution per se (e.g., the local Gini coefficient or a measure of a polarized, bimodal distribution) (see, e.g., Newman 2020; Newman and Kane 2017) as well as the presence of specific income groups vis-à-vis a given person's own income (e.g., a rich person's exposure to poor individuals) (see, e.g., Sands 2017; Sands and De Kadt 2020; Condon and Wichowsky 2020). The latter 'relational approach' can be viewed as an attempt to unpack which groups are driving a given inequality effect. While we follow the latter approach in this paper—our main operationalization of inequality is rich people's exposure to poor individuals—our study is informed by previous work based on both approaches.

The dominant finding in the literature is that residential exposure to higher local economic inequality in general—and exposure to poorer individuals among the well-off in particular—is associated with more support for redistribution and related attitudes (for example, Bailey et al. 2013; Franko and Livingston 2022; Kearns et al. 2014; Minkoff and Lyons 2019; Newman 2020; see also Franko 2016). These findings have been interpreted as being consistent with an extended version of intergroup contact theory, which predicts that contact with out-group members—particularly under facilitating conditions—can generate more positive attitudes toward these groups. This applies most straightforwardly to affluent individuals for whom the effect of such contact may extend beyond sympathy for poorer individuals to greater support for redistribution benefitting this group. While cross-income group contact is not the only explanation for the positive relationship, we label these studies 'contact-supporting' for simplicity.[2]

In sharp contrast to the general finding of a positive relationship between rich people's exposure to poorer individuals and their support for redistribution, Sands (2017) finds—in a prominent field experiment—that the presence of a noticeable poor-appearing person in affluent neighbourhoods in Boston, United States reduce passerby's redistribution support, gauged by their reduced willingness to sign a petition in support of a millionaire's tax. This finding is consistent with theories of intergroup conflict, which predict that exposure to economic out-groups prompts material conflict or evokes negative out-group stereotypes that, in turn, lead to a reduced willingness to redistribute among the better-off.[3,4] Correspondingly, we label Sands' finding 'conflict-supporting'.

The sharp divergence in their conclusions naturally raises the question of what may explain the difference between the contact- and conflict-supporting studies. At least two potential explanations lend themselves. *The self-selection explanation* attributes the diverging results to different abilities to eliminate bias from the selection of individuals into neighbourhoods in the two lines of work. Conversely, *the single episode vs. repeated exposure explanation* reconciles the divergent findings by attributing them to differences in the temporal extension of the out-group exposure studied as a proxy for associated intergroup processes (conflict or contact).

This paper proposes and tests the two potential explanations for the literature's divergent findings by analyzing the relationship between exposure to poor people and support for redistribution among rich individuals in Denmark, an economically equal country, comparatively speaking (Solt 2020). We do so by employing a design closer to the contact-supporting studies (focusing on repeated exposure to poor neighbours, providing the basis for eventual contact) but with more causal leverage than these (a three-wave individual-level panel design linked to fine-grained registry data on local income composition) to study the consequences of temporally extended residential exposure to poor people among the better-off. Our design thus allows us to address whether the divergent existing findings are (*i*) a function of the strength of causal identification in the design employed or (*ii*) the time horizon of the studied exposure.

To preview, in *within*-individual models analyzed using two-way fixed effects, we find that exposure to poor individuals is associated with *lower* support for redistribution among wealthier individuals. Our results resonate with Sands' (2017) findings in supporting the conflict perspective on the link between exposure to poor individuals and support for redistribution. Further, we substantiate *the self-selection explanation* by showing that when analyzed cross-sectionally (in *between*-individual models), we find a positive relationship between exposure to poor individuals

---

[2]For example, a similar prediction follows from Dimick, Rueda, and Stegmüller's (2018) theory of income-dependent altruism, although instead of increased sympathy for the poor specifically, this theory holds that rich individuals become more altruistic in general, when faced with inequality.

[3]Relatedly, in a recent study, Sands and de Kadt (2020) show that exposure to wealth prompts conflict-responses (i.e., higher support for redistribution) among poor individuals.

[4]At more aggregate subnational levels—e.g., American States—there is mixed evidence for the conflict perspective (see, e.g., Côtè, House, and Willer 2015; Franko 2016).

and redistribution support among the better-off, thus indicating that self-selection based on stable individual characteristics is a likely explanation for previous 'contact-supporting' findings. Our results thus support the negative effect of exposure to poor individuals on the better-off's support for redistribution and point to one plausible explanation for why some studies have found the opposite.

Self-Selection or Different Time Frames? Explaining the Divergent Results

As noted, two accounts—*the self-selection explanation* and *the single episode vs. repeated exposure explanation*—may plausibly explain the divergent findings in the literature on the relationship between local inequality and support for redistribution.

The self-selection explanation takes its starting point in the fact that contact-supporting studies rely on observational—almost exclusively cross-sectional—data, while Sands' (2017) conflict-supporting study is based on an experiment. Consequently, one explanation for the different findings is that the positive correlation between inequality and support for redistribution observed in the contact-supporting studies reflects confounding from self-selection. More specifically, rich individuals—who, given their financial resources, are less constrained in where they live—may partly *choose* neighbourhoods and neighbours based on their predispositions toward the less well-off. Because these predispositions likely correlate with attitudes toward redistribution, the contact-supporting evidence may reflect a choice to live in certain areas that is correlated with redistribution attitudes. Conversely, because Sands (2017) randomizes exposure to poor individuals, the observed relationship is free from selection bias. Lastly, if *the self-selection explanation* is correct, a logical follow-up question is which selection processes—for example, contemporary experiences or personality traits—explain the difference in the results.

The single episode vs. repeated exposure explanation attributes the literature's divergent results to the time horizon of the observed relationship between exposure to poor individuals and support for redistribution. Rather than siding with either line of work, this explanation attributes the divergent findings to the difference between short-term/single-shot (conflict-supporting studies) and repeated (contact-supporting studies) exposure. More specifically, theory and empirical evidence from other domains, especially immigration, suggest that it is exactly the fleeting exposure to out-groups, which Sands studies, that likely prompts adverse reactions to these groups (Allport 1954; Enos 2017). As such, and by extension, the reduction in support for redistribution among the better-off when briefly exposed to poorer individuals aligns perfectly with the broader literature on group conflict. However, whether repeated out-group exposure consolidates or moderates the negative effect so that it (possibly) becomes positive is a matter of debate. Inferring from related research—for example, on the relationship between local ethnic diversity and social trust (Dinesen and Sønderskov 2015)—the type of transient exposure studied by Sands might accumulate over time to manifest itself in lower support for redistribution more generally. Yet, consistent with the literature on out-group contact, exposure may also come to resemble 'actual contact' (that is, meaningful personal interactions) when occurring over extended periods of time (Allport 1954; Pettigrew and Tropp 2011). Consequently, such contact could give rise to the positive relationship between exposure to inequality and support for redistribution observed in the contact-supporting work, which arguably studies repeated exposure resulting from the general socio-economic composition in an area, which only changes gradually over time.[5] In short, the contrasting results from the two

---

[5]More specifically, existing contact-supporting studies, which link traditional surveys to administrative data, measure more prolonged or repeated exposure to poor people living nearby than the transient single-episode exposure that Sands' study. This type of exposure, occurring over extended periods, may come to resemble meaningful personal interactions (Allport 1954; Pettigrew and Tropp 2011), which may then in turn explain the divergent findings.

lines of work observing conflict and contact effects may potentially be reconcilable given the different time horizons they implicitly study.

Research Design, Data, and Measurement

Our study utilizes a within-individual research design linking panel survey data with fine-grained individual-level registry data. The panel data are from the Social and Political Panel Study, which is a three-wave panel collected in 2008/9 (*wave 1*, collected as part of the European Social Survey), 2011/12 (*wave 2*), and 2017 (*wave 3*). Appendix A.1 describes the panel data in detail.

The panel survey data are linked to individual-level registry data from Statistics Denmark, which contains longitudinal data on the exact geographical location of the place of residence for all individuals with primary residence in Denmark, except for some EU citizens, illegal residents, and asylum seekers (see Appendix A.2.1), and extensive socio-economic information, including personal income, reported by public authorities. This allows us to construct measures of the income composition of individualized local contexts flexibly as circles with a given radius around the respondents. The precise and flexible measures of local inequality (and related variables, including potentially confounding factors) provide for a strong test of the proposed relationship between local economic conditions and support for redistribution compared to studies measuring exposure in more aggregate contexts (for example, US counties) that may be imprecise reflections of people's everyday experiences of income differences (Bisgaard, Dinesen, and Sønderskov 2016).

Our dependent variable, support for redistribution, is measured by the widely used question included in the European Social Survey: '[the] Government should take measures to reduce differences in income levels' (for example, Dimick, Rueda, and Stegmueller 2018)—assessed on a five-point Likert-scale from 'Agree strongly' (1) to 'Disagree strongly' (5)' (reversed in the analysis). Focusing on a relatively general orientation rather than a stance on a specific policy (for example, the millionaire's tax in Sands' study) renders our test conservative, given the importance of 'linkability' between policies and attitudes (Citrin and Green 1990).

Inspired by previous studies (Condon and Wichowsky 2020; Franko and Livingston 2022), the independent variable in the primary analyses—exposure to poor individuals—is measured as the share of residents in a local context (of a given size) with an income below the 20th percentile in the national income distribution. In robustness checks, we employ two alternative measures of the independent variable (see Appendix B.1). To examine the local foundation of our results, we use measures of exposure to poor individuals in a range of context sizes, specifically, circles with radii from 100 meters (=0.031 km^2; our primary specification) up to 2,500 meters. The smallest, ultra-local context is our primary focus because the immediate local context is more likely to capture quotidian exposure to social phenomena (Bisgaard, Dinesen, and Sønderskov 2016; Danckert, Dinesen, and Sønderskov 2017; Dinesen and Sønderskov 2015)—in our case, poverty.[6]

Figures A8–A9 in the Appendix provide information about the correlates of within-individual changes in the share of poor individuals over time. All observed covariates were only weakly correlated with changes in exposure to poor people ($r < 0.17$). Furthermore, neither positive nor negative changes seem to be concentrated in specific geographic regions (Figure A10). We also examined whether changes in exposure to poor people are driven by replacement of residents (people moving in/out), changes in income of current residents, or relocation of respondents themselves in Appendix B.4.3 (note that we excluded observations of recently relocated

---

[6]As we discuss in Appendix A.1, we exclude person-wave observations of persons who, at the time of interview, had moved to their current address within 6 months prior to the interview as well as observations living in sparsely populated contexts (less than 15 neighbors and 2 families living within 100 meters) to increase validity and reliability of our contextual measures. In Appendix D.3, we examine the sensitivity of our results to these sample restrictions. The estimates based on models without restrictions are slightly attenuated but neither substantially nor significantly different from those based on our primary sample.

respondents altogether, see footnote 6). The analyses show that all three kinds of changes contribute to the variation in exposure to poor people. These analyses indicate that changes in exposure to poor individuals are not confined to particular places or processes.
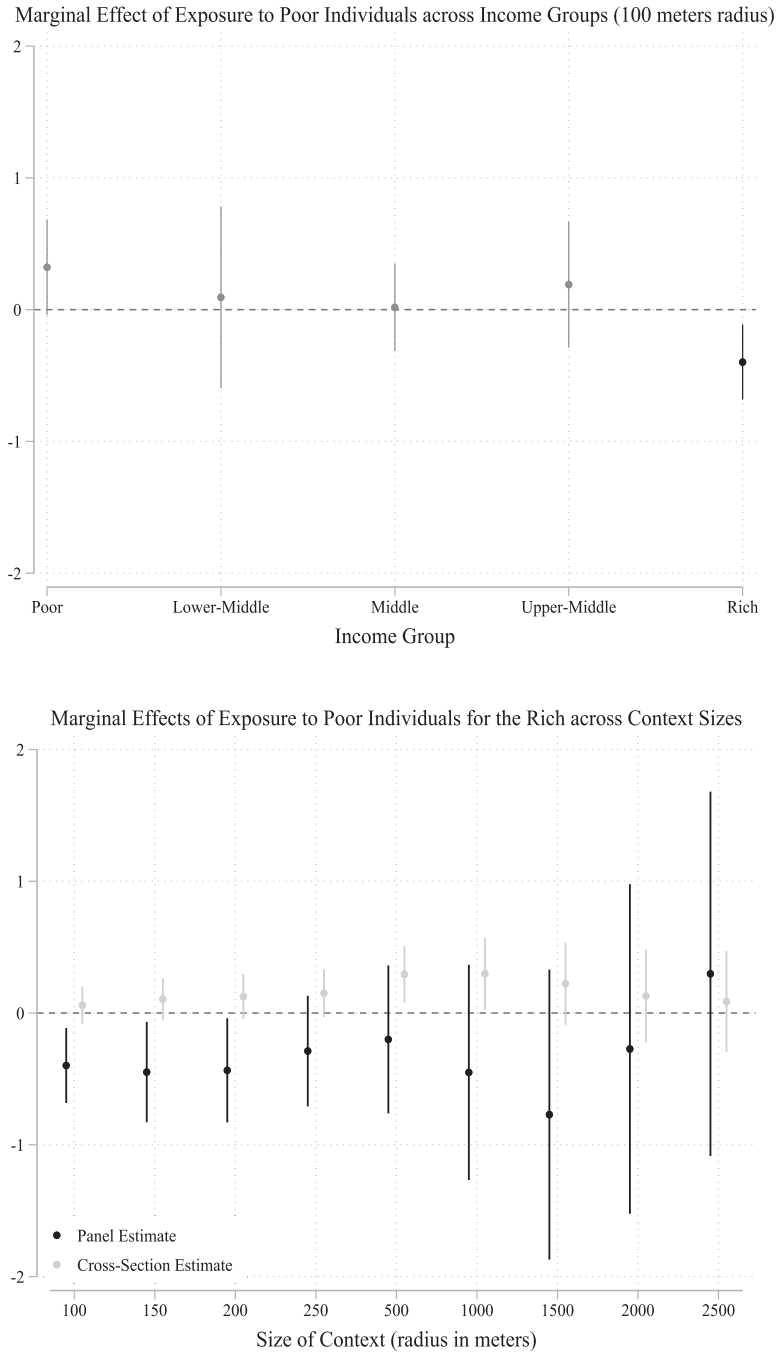
We analyzed the panel data by means of two-way—individual and time—fixed-effects (TWFE) models, which analyze the relationship between exposure to the poor and support for redistribution *within* individuals over time and net of any general temporal fluctuations in support for redistribution. This strategy allows us to rule out confounding of the relationship between exposure to poor people and redistribution support from time-invariant factors. In addition, we include a rich set of time-varying registry-based individual- and contextual-level control variables that potentially confound the relationship (see Table A1 in Appendix B.2). Parallel to related work (Danckert, Dinesen, and Sønderskov 2017), we do not include measures of political ideology (or other predispositions) out of concerns for post-treatment bias (Angrist and Pischke 2009) in the fixed effects model that only analyze over-time changes.[7] To analyze the relationship separately for income groups, we interact exposure to the poor with respondents' income quintile in the first survey wave in which they are included in the analyses. Lastly, to gauge potential time-invariant sources of confounding, we also analyzed the relationship between exposure to poor individuals and support for redistribution in cross-sectional models (that is, without individual fixed effects).

Results

The upper panel in Fig. 1 presents the marginal effect of the share of poor people in the local context across five income quintiles based on Model 1, Table A8 in the Appendix. As Fig. 1 illustrates, the marginal effect of the share of poor in the local context is, statistically, significantly negative (-0.40, $p < 0.01$) for the top income quintile. By contrast, the corresponding marginal effects are insignificant for all other income groups, and they differ significantly from the estimate for the richest group (except for the lower-middle income group). Further, the negative effect of exposure to poor individuals for the richest quintile is quite substantial; a one standard deviation increase in the local share of poor decreases support for redistribution by 5.1 percentage points. By comparison, Sands (2017) finds that subjects become 4.4 percentage points less supportive of a millionaire's tax when exposed to a noticeably poor-appearing person. Further, Fig. 1 (lower panel, dark lines) indicates that the estimate is only *significantly* negative in smaller contexts (up to a radius of a couple of hundred meters). This tentatively suggests a rather local nature of the effect of exposure to the poor on the well-off's redistribution attitudes (Sands and de Kadt 2020), but the imprecise estimates in larger contexts prevent us from drawing any firm conclusions about this. In short, our results demonstrate that exposure to poor individuals is associated with lower support for redistribution among the richest individuals, thus substantiating the conflict perspective and, by implication, speaking in favour of *the self-selection explanation* for the literature's diverging results.

How do we know that our results are driven by exposure to poor individuals, given that this group share is mechanically correlated with the share of other income groups? Or, put more substantively, perhaps it is the psychological comfort of being among more of their own kind (and thus fewer poor individuals) that makes richer individuals less sceptical of redistribution. Changes in the share of poor and rich individuals, respectively, are relatively highly correlated

---

[7]To the extent that these predispositions do change over time and therefore, by implication, are not accounted for by the fixed effects, their changes may themselves be a consequence of local exposure to poor individuals, which would introduce post-treatment bias—most likely a downward bias—in the estimated relationship between the latter and support for redistribution. To facilitate comparison with earlier work, we tried including a measure of political ideology (left-right self-placement) in the cross-sectional models (see footnote 10).

Marginal Effect of Exposure to Poor Individuals across Income Groups (100 meters radius)

Marginal Effects of Exposure to Poor Individuals for the Rich across Context Sizes

Figure 1. Marginal effects of exposure to poor individuals across income groups and context sizes.
Note: The upper panel is based on the estimates reported in Model 1 in Table A8. The panel estimates in the lower panel are based on specifications similar to Model 1 in Table A8 but with varying contextual sizes. The cross-sectional estimates are based on specifications similar to Model 1 in Table A9.

($r = -0.40$), making it difficult to distinguish their effects empirically. Nevertheless, when adding the share of rich and its interaction with income quintiles to our baseline specification, we observe that while the negative marginal effect of exposure to the poor among the richest quintile

is reduced (from −0.40 to −0.34), it remains relatively strong and statistically significant (see Appendix Table A8, Model 2). Relatedly, adding the local Gini coefficient (see Appendix B.2)—a commonly used indicator of inequality, especially in country-comparative work (for example, Dimick, Rueda, and Stegmueller 2018)—and its interaction with income quintiles (Model 3, Table A8), also leaves the estimated association for the richest individuals essentially unchanged.[8] Similarly, repeating this analysis with a measure capturing income diversity (Minkoff and Lyons 2019) does not affect the results (Model 4, Table A8). Surviving these demanding tests speaks to the robustness of the relationship, and it carries the important theoretical implication that it is specifically local exposure to poor individuals that reduces the well-off's support for redistribution, not living among fewer people from their own income bracket or an unequal income distribution in a more general sense.

We can further assess the self-selection explanation for the divergent results in the literature by comparing the estimated relationship between exposure to poor individuals and redistribution support among the better-off in cross-sectional and panel models, respectively. The dark (panel results) and light-grey (cross-sectional results) lines in the lower panel of Fig. 1 provide this comparison for different context sizes. Interestingly, while the relationship is consistently *negative* for the richest quintile in the panel models, it is—in line with results from previous studies using cross-sectional designs[9]—consistently *positive* (and sometimes statistically significant) in the cross-sectional models.[10] This striking divergence indicates that the selection, which plausibly accounts for the divergence between the existing conflict and contact-supporting studies, is rooted in time-invariant characteristics that are held constant in the panel models but not in the cross-sectional models. Further supporting the selection account for rich individuals is the fact that we find no such inversion in effects between the panel and the cross-sectional models for those in the lowest income quintile, who—compared to the well-off—are likely to be much more constrained in their choice of residence (see Appendix C.2). The results also hint at the sources of this time-invariant confounding. Given that the cross-sectional models control for a long list of precisely measured (registry-based) socio-demographic confounders, this suggests that the time-invariant confounding is not rooted in social deprivation broadly conceived. Instead, 'deep' psychological predispositions (for example, personality traits or values)—that cannot be measured through the registry data—appear as plausible time-invariant confounders. Put more substantively, when we, at a given point in time, observe that richer individuals are more supportive of redistribution when living in a neighbourhood with a higher share of poorer individuals, this reflects a sorting—plausibly based on deep-held predispositions—of more pro-redistribution individuals into these neighbourhoods rather than an effect of living among poor individuals.

We probe the negative effect of the share of poor individuals on redistribution support among the well-off in several robustness checks (reported in Appendices D.1–D.4). First, we show that our results are robust when applying a matching TWFE-estimator, where individuals experiencing positive and negative changes, respectively, are matched on pre-treatment characteristics to allow changes in redistribution attitudes to vary with these characteristics (see, for example,

---

[8]We also examined the unconditional and conditional effect of Gini on support for redistribution without the parallel interaction between the income quintile and the share of the poor. As reported in Table A10 in Appendix C.1, we find a negative but insignificant effect.

[9]Ceteris paribus, this also speaks in favour of our findings being representative of the larger universe of studies.

[10]In Figure A24 in the appendix, we report results from cross-sectional models that include left-right self-placement as a measure of ideology to facilitate comparison with earlier work including similar measures. As is evident from the figure, the cross-sectional estimates from models with and without left-right self-placement are virtually identical. Therefore, its inclusion has no implications for our inference regarding self-selection as a likely explanation for the divergent results in the existing literature.

Sant'Anna and Zhao 2020).[11] Second, our results are generally robust to using alternative measures of exposure to poor individuals (a measure based on half of the median income and a weighted (gradual) measure). Third, the results are also robust to various sample modifications; for example, excluding individuals younger than thirty years of age when calculating the share of poor individuals to avoid conflating poverty with age composition, limiting the sample to people staying put in the neighbourhood to (partly) address the consequences of self-selection in and out of neighbourhoods over time, and to drop restrictions on the length of residence as well as the number of people living within 100 meters of the respondents' residences. Fourth, the within-individual estimates of exposure to poor individuals remain negative and statistically significant when restricting the sample to include rich individuals only, thereby addressing income-group specific confounding from the control variables. Fifth, we find that past changes in support for redistribution do not significantly predict subsequent changes in the local share of poor individuals, thereby supporting the parallel trends assumption. Although insignificant, the association between prior changes in redistribution attitudes and subsequent exposure is positive, tentatively pointing to the self-selection explanation as a plausible account for the positive relationship found in contact-supporting studies. Collectively, these analyses suggest that our results are rather robust.

We also conducted additional analyses to explore the conditionalities and scope of our results (see Appendix D.5). First, we find that the effect is primarily driven by positive changes in exposure to poor individuals, suggesting that the rich respond specifically to *increases* in exposure to poor individuals. Second, consistent with the literature on politically motivated information processing (Bisgaard 2015; Taber and Lodge 2006), political predispositions may condition responses to exposure to poor individuals. More specifically, in the case of political ideology, exposure to a higher share of poor people locally might trigger stronger opposition to redistribution among ideologically right-leaning, well-off individuals. We examine this conjecture by interacting political ideology (measured in the first wave of the panel survey) with subsequent exposure to poor individuals among the rich (see Table A19 and Figure A23, Panel A in the appendix). We find no support for ideologically conditioned responses to local exposure to poor individuals. Third, we tried using political ideology as the dependent variable to examine if the effect of exposure to poor individuals extends from a specific aspect of distributional politics to more general political orientations among the rich. While in a similar direction as the results for support for redistribution, the estimated effects for political ideology are substantially weaker and statistically insignificant (see Model 2, Table A19 in the Appendix). Thus, living among poor neighbors only to a limited extent manifests itself in rich individuals' broader political ideology. Fourth, a placebo-like test shows that exposure to poor individuals does not matter with regard to immigration attitudes, indicating that the effect is restricted to classical distributional politics, not newer, more value-based issues.

Lastly, while not our primary focus, it is worth noting that higher local exposure to poor individuals generally significantly *increases* support for redistribution among the lowest income quintile. Conversely, and in contrast to Sands and de Kadt's (2020) findings from South Africa, we find no systematic effect of exposure to a higher share of rich individuals for this income group (see Appendix E). One explanation for this result might be that living with more people of the same low economic standing furthers a sense of group solidarity. More importantly, when seen in conjunction with the primary result, this suggests that it is specifically exposure to poor individuals that is consequential for redistribution attitudes—*both* for the bottom and the top income quintile. We conjecture that this might have to do with the visibility of this group, but this remains speculation and merits further attention in future research.

---

[11]We also estimated the effect for each period to accommodate potential negative weighting bias that may affect the ATT estimate in panel models with more than one time period (Chaisemartin and D'Haultfoeuille 2022). Again, we find a negative effect among the well-off, although imprecisely estimated in the first period (see Figure A13).

Conclusion and Discussion

This paper aims to explain the divergent findings in the literature on the relationship between local inequality—specifically exposure to poor individuals among the well-off—and support for redistribution. In short, we find that local exposure to poor individuals reduces support for redistribution among the well-off—supporting the conflict perspective—and that self-selection is a likely explanation for why previous studies have found that exposure to poor individuals furthers support for redistribution.

Our study contributes by expanding the scope conditions of the negative effect of exposure to poor individuals on support for redistribution among the better-off observed in Sands' (2017) foundational study. First, we show that the negative effect of exposure to poor individuals ostensibly consolidates over time and thus extends beyond a transient effect after a specific episode. Second, beyond the millionaire tax studied by Sands (2017), we show that the negative effect is observable in a more abstract outcome by tapping general attitudes toward redistribution. Third, our findings show that the negative effect of exposure to poorer individuals for redistribution support extends to a relatively economically equal Scandinavian welfare state where income differences are likely to be less pronounced and visible. In short, our findings significantly extend the scope of the negative effect of exposure to poor individuals on support for redistribution among the better-off.

At the same time, our work has limitations and raises new questions. Theoretically, we still have only a relatively rudimentary understanding of why exposure to poor individuals makes rich people less supportive of redistribution. Exploring what exact mechanisms—for example, stereotype activation or negative out-group affect—underlie this 'conflict response' and why we do not see a symmetrical response when poor people are exposed to rich individuals are relevant foci for future research. Further, while we have shown that self-selection might account for the 'contact-supporting' findings in the literature, discarding the extended version of contact theory based on our results would be mistaken. Like others, we have only addressed intergroup contact by proxy of temporal exposure (arguably 'contact potential'). This is clearly not an exhaustive measure of contact—indeed, there are situations where contact may never emerge, even with temporally extended exposure—and there are other 'facilitating conditions'—friendship potential, common goals, etc.—that have been shown to promote positive intergroup relations (Allport 1954; Pettigrew and Tropp 2011). Further, work on the consequences of ethnic diversity has shown that contact can moderate conflict responses (Stolle, Soroka, and Johnston 2008), thereby pointing to the potential interplay between conflict and contact. In this regard, one avenue for further research would be to study the consequences of economic out-group exposure in other, more 'contact-prone' contexts than the residential one. Building on recent research (Rao 2019), schools—where parents work for a common purpose—would be an interesting setting in this regard.

Lastly, finding that local exposure to poor individuals decreases support for redistribution among the well-off highlights the challenges associated with ameliorating high levels of economic inequality. Rather than animating a self-correcting demand for greater equality, such experiences tend to perpetuate inequality by lowering the desire for redistribution among those who have the means to redistribute. At the same time, our finding that poor individuals respond to local exposure to other poor people by becoming more supportive of redistribution provides a silver lining for those hoping to redress existing inequality. Ultimately, the contrasting reactions to local poverty between the rich and the poor highlight the material conditions underlying political conflicts around inequality.

References

Allport GW (1954) *The Nature of Prejudice*. Reading, MA: Addison-Wesley.

Angrist JD and Pischke J (2009) *Mostly Harmless Econometrics*. Princeton: Princeton University Press.

Bailey N, Gannon M, Kearns A, Livingston M and Leyland A (2013) Living apart, losing sympathy? How neighbourhood context affects attitudes to redistribution and to welfare recipients. *Environment and Planning* **45**, 2154–75.

Bisgaard M (2015) Bias will find a way. *Journal of Politics* **77**(3), 849–60.

Bisgaard M, Dinesen PT and Sønderskov KM (2016) Reconsidering the neighborhood effect: Does exposure to residential unemployment influence voters' perceptions of the national economy? *The Journal of Politics* **78**(3), 719–32.

Chaisemartin C and D'Haultfoeuille X (2022) TWFE And differences-in-differences with heterogeneous treatment effects. *Econometrics Journal* **236**(2), 1–20.

Christensen ME, Dinesen PT and Sønderskov KM (2024) Replication Dofile for: Unequal and Unsupportive: Exposure to Poor People Weakens Support for Redistribution among the Rich. Available from https://doi.org/10.7910/DVN/GCCULO, Havard Dataverse, DRAFT VERSION.

Citrin J and Green DP (1990) The self-interest motive in American public opinion. *Research in Micropolitics* **3**(1), 1–28.

Condon M and Wichowsky A (2020) *The Economic Other*. Chicago: University of Chicago Press.

Côte S, House J and Willer R (2015) High economic inequality leads higher-income individuals to be less generous. *Proceedings of the National Academy of Sciences* **112**(52), 15,838–843.

Danckert B, Dinesen P and Sønderskov KM (2017) Reacting to neighborhood cues? *Public Opinion Quarterly* **81**(1), 37–56.

Dimick M, Rueda D and Stegmueller D (2018) Models of other-regarding preferences, inequality, and redistribution. *Annual Review of Political Science* **21**(1), 441–60.

Dinesen PT and Sønderskov KM (2015) Ethnic diversity and social trust. *American Sociological Review* **80**(3), 550–73.

Enos R (2017) *The Space Between Us*. New York: Cambridge University Press.

Franko WW (2016) Political context, government redistribution, and the public's response to growing economic inequality. *The Journal of Politics* **78**(4), 957–73.

Franko WW and Livingston AC (2022) Economic segregation and public support for redistribution. *The Social Science Journal* **59**(3), 1–19.

Kearns A, Bailey N, Gannon M, Livingston M and Leyland A (2014) All in it together? *Journal of Social Policy* **43**(3), 453–77.

Kelly NJ and Enns PK (2010) Inequality and the dynamics of public opinion. *American Journal of Political Science* **54**(4), 855–70.

Kenworthy L and McCall L (2008) Inequality, public opinion, and redistribution. *Socio-Economic Review* **6**(1), 35–68.

Lupu N and Pontusson J (2011) The structure of inequality and the politics of redistribution. *American Political Science Review* **105**(2), 316–36.

Meltzer AH and Richard SF (1981) A rational theory of the size of government. *Journal of Political Economy* **89**(4), 914–27.

Minkoff SL and Lyons J (2019) Living with inequality. *American Politics Research* **47**(2), 329–61.

Newman BJ (2020) Inequality growth and economic policy liberalism. *The Journal of Politics* **82**(2), 765–70.

Newman BJ and Kane JV (2017) Economic inequality and public support for organized labor. *Political Research Quarterly* **70**(4), 917–32.

Pettigrew T and Tropp L (2011) *When Groups Meet. The Dynamics of Intergroup Contact*. New York: Psychology Press.

Piketty T and Saez E (2014) Inequality in the long run. *Science (New York, N.Y.)* **344**(6186), 838–43.

Rao G (2019) Familiarity does not breed contempt. *American Economic Review* **109**(3), 774–809.

Sands ML (2017) Exposure to inequality affects support for redistribution. *Proceedings of the National Academy of Sciences* **114**(4), 663–8.

Sands M and de Kadt D (2020) Local exposure to inequality raises support of people of low wealth for taxing the wealthy. *Nature* **586**(7828), 257–61.

Sant'Anna PHC and Zhao J (2020) Doubly robust difference-in-difference estimators. *Journal of Econometrics* **1**(219), 101–22.

Solt F (2020) Measuring income inequality across countries and over time. *Social Science Quarterly* **101**(3), 1183–99.

Stolle D, Soroka S and Johnston R (2008) When does diversity erode trust? *Political Studies* **56**(1), 57–75.

Taber CS and Lodge M (2006) Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* **50**(3), 755–69.