Neighbour Noise, Well-being and Health: Mixed-Methods Evidence from Chile

Àlex Boso, Diana Weinhold* & Jaime Garrido

August 23, 2026

Abstract

Neighbour noise is common, difficult to observe before choosing a dwelling, and unusually entangled with social relations and the permeability of the home. Yet most empirical evidence on its consequences is subject to confounding biases that make it difficult to distinguish an effect of noise from residential selection or outcome-driven reporting. We combine seven waves of the Chilean Longitudinal Social Survey with 45 interview sessions conducted in Temuco, Chile, comparing between- and within-person estimates, examining selection into exposure, and the direction of timing and sensitivity to omitted variables. We then triangulate our quantitative results with extensive qualitative evidence that allows us to better understand the social context of noise exposure and build confidence in the interpretation of observed patterns. In the quantitative data the same respondent reports lower life satisfaction, more sleep difficulty and more depressive symptoms in noisier waves. These within-person relationships are smaller than the between-person differences, but remain statistically significant and change very little as time-varying socioeconomic circumstances and residential history are added. Prior outcomes do not consistently predict the onset of high reported noise. The interviews give these patterns substantive content, highlighting the importance of recurring night-time disturbance, disrupted sleep, anticipatory stress, social meaning and limited possibilities for obtaining relief. While confounding is still possible, the longitudinal stability, selection checks and cross-method concordance make a wholly compositional or outcome-generated reporting account difficult to sustain.

*Corresponding author: Department of International Development, London School of Economics and Political Science, London, UK; d.weinhold@lse.ac.uk.

1 Introduction

As many an urban dweller can attest, neighbour noise is a particularly vexing aspect of life in the city. Of the 3.43 million requests New Yorkers made to the city's '311' non-emergency line in 2024, the largest single category (18 percent) was noise, and 62 percent of these were for residential noise—loud music, parties, neighbours. Complaints about the people next door thus accounted for eleven percent of everything New Yorkers asked of their municipal government, more than heat and hot water, blocked driveways or the condition of the streets (Office of the New York State Comptroller 2025).

Neighbour noise is a reciprocal externality; as has been understood since Coase (1960), the efficient outcome depends on how rights are assigned and on what it costs to bargain over them. In the case of residential noise, rights can be ambiguous, even innocent noise can generate social conflict, and bargaining over noise is intrinsically socially expensive (Stokoe & Hepburn 2005). Transmission of neighbour noise is thus a disamenity for both the producers and the hearers, and because neighbour noise is not generally observable *ex ante*, its presence does not reliably get capitalised into rents and prices the way a busy road does, and the shadow price of a loud (or sensitive) neighbour manifests instead in ways less easy to measure—in life satisfaction, health, and relocation costs (Weinhold 2013, 2014). Praag & Baarsma (2005), for example, found that the disamenity of Amsterdam airport noise was absorbed not in house prices but as a residual in the life satisfaction of owners who could not move.

If neighbour noise is difficult to ameliorate privately, this points to an important role for public policy. Building standards have generally been improving for modern construction, but much of the housing stock in many cities still provides inadequate insulation, as the 311 calls can attest, and local regulations are uneven; the WHO's 2018 guidelines for the European Region, which set exposure limits for road, rail and air traffic, wind turbines and leisure noise, say nothing about domestic transmission of neighbour noise, either in the guidelines themselves or in the systematic reviews commissioned to support them (World Health Organization Regional Office for Europe 2018, Jensen et al. 2019). Where neighbour noise is attenuated not by insulation but by rules and norms, enforcement generally falls under local nuisance law rather than an environmental health authority, and is left to local governments with widely varying priorities and capacity (Hammer et al. 2018). In New York City itself, from 2010 to 2015 the NYPD investigated 1.3 million noise complaints, but only 0.4 percent of these turned into formal summonses (Office of the New York State Comptroller 2018). Chile provides a typical case in point. The national noise standard, D.S. 38/2011, sets maximum permissible levels for regulated fixed sources—55 dB by day and 45 dB by night in residential zones—and then expressly exempts the activity proper to inhabiting a dwelling: voices, gatherings, pets, appliances, domestic repairs (Ministerio del Medio Ambiente, Gobierno de Chile 2011). What remains is handled by municipal ordinance, by the internal co-ownership rules condominiums must maintain under Law 21.442 (Congreso Nacional de Chile 2022), and, when a night turns disorderly, by *Carabineros*, the local police.

The lack of consistent and effective public policy on neighbour noise in turn reflects a dearth of robust evidence of harm upon which to base policy action. The literature on traffic and aircraft noise is large because those exposures can be objectively measured with acoustical meters; decibels can be modelled from emissions and mapped over space, and large prospective cohorts or quasi-experimental variation allow physical dose–response functions to be estimated (Babisch 2014, Sørensen et al. 2013, Smith et al. 2022, Nissenbaum et al. 2012, Boes et al. 2013, Wang et al. 2022, Beghelli et al. 2023). Neighbour noise differs from traffic, air and other environmental sources of noise in that it arrives with a social context: voices, footsteps, music and arguments carry information, and human attention is drawn to informative sound even

when it is quiet (Niemann et al. 2006). The biological mechanisms that run through sleep disruption and stress respond not only to sound level but to timing, recurrence, and social context—properties the interviews in Section 4 show respondents distinguishing carefully but which are very hard to consistently measure. Instead, data on the incidence of neighbour noise annoyance is almost exclusively collected as subjective, self-reported binary or scaled responses in large, multi-purpose household surveys (e.g. Niemann et al. 2006, Weinhold 2015, Jensen et al. 2019). In this data, variation in self-reported neighbour noise does not arise experimentally but is *endogenous* — any measured correlation with health and well-being outcomes could arise either from a causal effect of noise, or spuriously from a confounding correlation. Households with the fewest housing options may end up in the dwellings that transmit the most sound, so that the correlation simply records who lives where. Broader neighbourhood deterioration may raise noise and lower well-being together, with no direct link running between them. And poor sleep, stress and depression may themselves change what a resident notices and reports, so that the causation runs from lower well-being to reported noise rather than the reverse. Adding controls can mitigate, but does not fully solve, these methodological problems; confounders can vary over time and many are plausible consequences themselves of an unresolved noise problem.

If robust causal empirical evidence of the harm from neighbour noise is still elusive, the theoretical case has been strengthening. The medical literature has increasingly documented confirmed biological pathways through which stress and noise exposure induces inflammation, especially in vascular and cerebral tissue (Osborne et al. 2020, Münzel et al. 2021, 2024, Kröller-Schön et al. 2018), so there is a sound theoretical basis to suspect that exposure to neighbour noise could have significant effects on mental and physical health. And, despite the methodological difficulties, suggestive empirical evidence of significant ill effects from neighbour noise has been accumulating. In a nationwide Danish sample of adults living in multi-storey housing, those very annoyed by neighbour noise had between 1.7 and 3.3 times the odds of being severely bothered by each of eight physical and mental symptoms, from headache and fatigue to musculoskeletal pain, depression and anxiety (Jensen et al. 2019). The pan-European LARES study found elevated risks across the cardiovascular, respiratory and musculoskeletal systems as well as depression, with a dose–response gradient that survived adjustment for socioeconomic and housing conditions (Niemann & Maschke 2004, Niemann et al. 2006). Earlier work on Dutch panel data found comparable magnitudes: the odds of cardiovascular disease among those reporting noisy neighbours were 1.38 times those of respondents with quieter ones, and the odds of fatigue nearly four times (Weinhold 2015).

Longitudinal survey evidence, which is able to control for time-invariant factors and separate noise and effects over time, is rarer, but where it exists has been especially useful. Analysis over time of initially healthy individual respondents of the Dutch LISS panel finds associations between neighbour noise and cardiovascular symptoms, joint and bone disease and headache (Weinhold 2015). In the Gutenberg Health Study, baseline noise annoyance predicted new depression, anxiety and sleep disturbance five years later (Beutel et al. 2020). Using eight waves of the same Dutch panel, Cassens et al. (2025) identify chronic, occasional and never-annoyed profiles, find roughly twelve percent of respondents chronically annoyed by neighbours, and estimate cross-lagged models in which chronic annoyance raises sleep problems while deteriorating heart complaints also raise subsequent annoyance.

We contribute to this growing body of evidence by combining quantitative analysis from a novel (to this literature) longitudinal survey with extensive qualitative interviews that show how social context moderates the reaction to noise. The quantitative analysis uses seven waves of the Chilean Longitudinal Social Survey (ELSOC), a nationally representative panel of the urban population, which allows us to compare people with different average exposure, to compare the same person in waves when she reports more and less neighbour noise, and to ask whether

prior outcomes and prior residential circumstances predict who begins reporting noise. The qualitative analysis, on a separate sample from the survey data, draws on 45 interview and household sessions in Temuco, Chile, that establish the source of the sound that generated a reaction, its hour, how often it recurred, the relationship with the neighbour who produced it, and what happened when the respondent tried to stop it. Chile also supplies a setting the literature has largely missed; evidence from middle-income countries is thin and geographically concentrated (Chen et al. 2023), and the norms governing what neighbours owe one another in a Latin American city need not resemble those of the northern European populations from which nearly all existing estimates come.

The longitudinal data allow us to separate respondents reporting persistently high noise annoyance from those reporting persistently low annoyance throughout the period and those whose annoyance changes. Of the 1,911 respondents observed in every wave from 2016 to 2019, 6.1 percent report high noise throughout, 36.0 percent remain persistently low, and 58.0 percent cross the threshold at least once. Several striking patterns stand out. First, conventional socioeconomic characteristics barely distinguish the persistently high, changing, and persistently low groups: education, university completion, household income, unemployment and financial precarity are close to identical across the three categories. Second, prior well-being, sleep and depressive symptoms do not predict the onset of a high-noise report, and next-wave noise is unrelated to the current outcome. Third, residential appraisal moves together with noise reports: relative to the persistently low group, the persistently exposed report neighbourhood safety 0.96 standard deviations lower and neighbour trust 0.84 standard deviations lower, and are some 19 to 20 percentage points more likely to want to move.

Consistent with previous studies, we also find strikingly robust correlations with well-being and health outcomes, both in the cross section and within individual respondents over time. Comparing respondents with different average exposure, a one-category increase in average reported neighbour noise is associated with 0.11 points lower life satisfaction, 0.19 points more sleep difficulty and 0.11 points higher mean depressive symptoms, each measured on a five-point scale. Among the changers, in waves when the same respondent reports more neighbour noise she also reports lower life satisfaction, more sleep difficulty, and more depressive symptoms, all with $p < 0.001$. The relationships remain robust to controlling for respondent fixed effects, contemporaneous changes in socio-economic variables, exclusion of the pandemic-era telephone wave, alternative binary thresholds and a fully non-linear exposure specification.

These within-person magnitudes are modest in standard-deviation terms but comparable to familiar life events. Moving from low to high reported noise is associated with a deterioration in life satisfaction that almost exactly offsets the improvement associated with entering a partnership, and with an increase in depressive symptoms slightly larger than that associated with becoming unemployed. We also report an association between persistent exposure and incident heart disease among respondents aged 50 and over which, albeit on limited data, reproduces the pattern found in the Dutch data a decade earlier, including the absence of any hypertension counterpart on a larger risk set.

The qualitative interviews show how noise is interpreted and translated into stress and annoyance. This evidence suggests that noise reactions are especially sensitive to social context, recurrence, and night-time occurrence. Eighty percent of neighbour annoyance episodes are recurring patterns, night-time disturbance appears in 35 of 43 exposed sessions, and sleep disruption and hypervigilance each appear in more than half. Respondents describe remedies that quiet a neighbour for a week, vigilance that begins before the sound does, and relief when a source moves out. They also describe habituation, sounds they consider entirely reasonable, and periods when competing demands such as studying made ordinary noise harder to tolerate—23

of the 43 exposed sessions contain both an adverse and a benign appraisal.

The paper proceeds as follows. Section 2 develops a conceptual framework in which transmission determines what enters the home, appraisal shapes how the sound is experienced, and control determines whether the disturbance can be resolved or becomes recurrent. Section 3 describes the quantitative methodology and presents the panel evidence on between- and within-person relationships, selection and timing. Section 4 presents the qualitative approach and examines the sources, meanings and responses represented by reported disturbance. Section 5 provides a triangulating synthesis of the combined bodies of evidence. A concluding discussion in Section 6 returns to interpretation, limitations and policy. An Appendix provides supporting material; Appendix Appendix A presents the extended battery of quantitative robustness checks, and Appendix Appendix B provides details of the structured qualitative coding, construction and audit trail.

2 Conceptual framework: transmission, appraisal and control

Neighbour noise is not a single exposure that can be represented by sound pressure alone. *Neighbour sound* is the physical acoustic event generated in or around another dwelling. *Experienced neighbour-noise disturbance* is the residential stressor that arises when that sound crosses a dwelling boundary and is registered as intrusive, disruptive or threatening. *Reported disturbance* is the survey or interview measure of that experience. These components overlap, but they are not interchangeable—transmission depends on the source and on the acoustic permeability of the dwelling; experienced disturbance depends additionally on timing, recurrence, social meaning and perceived control; and reported disturbance may also reflect recall, scale use, household aggregation and the respondent's current state.

Stallen (1999), drawing on the stress and coping framework of Lazarus & Folkman (1984), models noise annoyance as the outcome of an appraisal process in which a listener evaluates what produced a sound, what it signifies, whether it was avoidable and what resources are available for dealing with it. Appraisal is not a substitute for acoustics; it is one of the processes through which an acoustic event acquires psychological and welfare salience. We thus organise the framework around three linked processes. *Transmission* determines what enters the home. *Appraisal* determines how the sound is interpreted and whether it becomes disturbing. *Control* determines whether the disturbance can be stopped, adapted to or escaped. The processes are analytically distinct but recursive rather than strictly sequential. Limited control can intensify appraisal; repeated disturbance can create anticipation before the next event; and conflict can alter the meaning of sounds that were initially ordinary. This framework explains why acoustic measures and reports of experienced disturbance answer different questions, and why the panel and interview evidence have complementary roles.

2.1 Sleep and stress pathways

Sound can directly delay sleep, interrupt it or fragment it, while an intrusive or threatening sound can activate a stress response even when it does not wake the listener (Schmidt et al. 2013). The two pathways may reinforce one another: repeated waking can increase next-day irritability and emotional vulnerability, while anticipation and monitoring can make it harder to fall asleep before any new disturbance occurs. Night-time exposure is therefore particularly important because it combines reduced background sound, an expectation of quiet and the possibility

of disrupting sleep (Smith et al. 2022, Münzel et al. 2024). Medical research has linked sleep disruption to disorders of the immune, metabolic, inflammatory, and mood regulation systems (e.g. Gangwisch 2014, Morgan & Tsai 2015, Shukla & Basheer 2016) and observational studies find strong correlations between sleep disruption and a variety of health problems, including cardiovascular disease, obesity, diabetes, depression, and respiratory illness (e.g. Covassin & Singh 2016, Palmer & Alfano 2016, Knutson et al. 2006, Zee & Turek 2006). Using longitudinal survey data from the Netherlands, Fan & Weinhold (2022) show that neighbour noise-moderated sleep disruption is associated with cardiovascular and auto-immune diseases as well as headache.

The physiological literature also provides a plausible account of how repeated stress could extend beyond immediate discomfort. Experimental and observational studies of transport noise associate exposure with activation of the hypothalamic–pituitary–adrenal axis and sympathetic nervous system, oxidative stress, endothelial dysfunction and inflammatory activity (Kröller-Schön et al. 2018, Münzel et al. 2021, 2024), while imaging and cohort evidence links modelled residential noise to stress-related neural activity, arterial inflammation and subsequent cardiovascular events (Osborne et al. 2020). Night-time aircraft noise has been shown to impair endothelial function (Schmidt et al. 2013), and chronic stress is also a proposed pathway to depression and anxiety, with the magnitude of the response depending partly on how an event is appraised and whether it continues to occupy attention (Segerstrom & Miller 2004, Hahad et al. 2025).

Experimental work also suggests that repeated transport-noise exposure need not produce complete cardiovascular habituation, and that recovery after exposure ends may differ across physiological systems (Frenis et al. 2021, Bayo Jimenez et al. 2023). While it remains an empirical question whether neighbour noise may display a similar pattern, it does suggest that not all self-described habituation should be automatically considered neutralizing.

2.2 Transmission and social appraisal

Before a sound can be appraised, it must enter the home. Shared walls, ceilings, floors, ventilation systems and structural vibration determine how much ordinary activity becomes audible across a dwelling boundary. Poor acoustic separation can produce what Walsh et al. (2021) call *obtrusive intimacy*: residents involuntarily learn the routines, arguments and schedules of people with whom they did not choose to share domestic space. In this setting privacy is not guaranteed by the physical boundary of the home but must be sustained collectively by construction, convention and restraint (Chernysheva 2024, Ureta 2007). Evidence that reported neighbour-noise complaints fall sharply under progressively stricter floor-impact-sound standards is consistent with transmission being a substantively important and policy-relevant part of the problem (Shin et al. 2026).

Once audible, neighbour sound also carries information that an impersonal source generally does not. Human attention is drawn to informative sound even when it is not especially loud (Niemann et al. 2006). A neighbour is a socially legible actor whose intentions may be inferred, whose routines become known and beside whom the listener expects to continue living. Accounts of neighbour noise therefore turn repeatedly on whether audibility represents ordinary intimacy or an intrusion and whether the source is viewed as friendly or hostile (Şentop Dümen & Rasmussen 2025*a*). Repetition can also produce learned aversion, particularly where low-frequency or building-borne sound is difficult to locate or escape (Leventhall 2004).

Timing, recurrence, predictability and perceived discretion consequently matter alongside physical intensity. A sound associated with necessary daytime activity may be tolerated even when it is conspicuous, while a sound understood as avoidable, performative or imposed may

be disturbing at a lower level. An open-ended disturbance can sustain monitoring because the listener does not know when it will stop. A recurring disturbance can also create a burden before the acoustic event begins, as the anticipated hour approaches. Thus physical transmission and appraisal jointly produce experienced disturbance.

2.3 Control vs. adaptation

Control concerns the capacity to alter either the source or its consequences. It may take the form of a direct request, recourse to a landlord, building administrator, municipality or police, physical modification of the dwelling, private coping or residential exit. A source that can be stopped may be less threatening than one expected to return after every complaint, while repeated failure to control can make each subsequent event signal both renewed disturbance and the absence of an effective remedy.

Neighbour noise is especially difficult in this respect because it is a reciprocal externality. Residents are simultaneously potential hearers and producers, the boundary between ordinary domestic activity and unreasonable disturbance is often contested, and direct bargaining can itself damage the relationship on which a quiet outcome depends (Stokoe & Hepburn 2005). Moving requires resources and creates relocation costs; in many cases it is not a feasible option. Adaptation must therefore be distinguished from resolution; habituation, coping, resignation and removal of the source are not equivalent outcomes.

This distinction also changes how the distribution of neighbour noise should be conceptualised. Research on environmental noise frequently finds greater exposure among lower-income and marginalised populations, although the evidence is geographically concentrated and dominated by transport sources (Trudeau et al. 2023). Transport noise is attached to identifiable infrastructure and can be capitalised into land values. Neighbour noise is often harder to observe before choosing a dwelling and can change when occupants or their routines change, so conventional income sorting may map less directly onto its incidence. That does not make the burden socially neutral, however: inequality may arise through the acoustic quality of the dwelling, the safety and social cost of confronting a source, access to an institution that will act, or the capacity to insulate or leave (Ureta 2007, Méndez & Otero 2018).

It is therefore useful to distinguish inequality in the probability of encountering a source from inequality in transmission and inequality in the capacity to obtain relief. A household need not be unusually likely to have a noisy neighbour in order to bear a disproportionate cost once disturbance begins. The empirical analysis consequently examines conventional socioeconomic sorting into reported disturbance separately from heterogeneity in its consequences and from residential constraint.

2.4 Measurement and implications for the empirical research design

As the previous discussion has emphasized, objective acoustic measures and subjective reports are not rival measurements of a single underlying variable. Acoustic instruments characterise the sound that reaches a location: its level, frequency, duration and timing. A report of experienced disturbance records whether the sound entered daily life as an intrusion or disruption, and for questions about welfare, sleep and stress, that experience is substantively relevant. The ELSOC survey we use in the quantitative analysis asks how often the respondent or another household member was annoyed or disturbed by noise arising as a problem with neighbours. We

therefore refer to it, where precision matters, as *reported experienced neighbour-noise disturbance*, abbreviated to *reported disturbance*, rather than as a direct measure of acoustic exposure.

Within this framework there a number of testable empirical implications. First, if reported disturbance captures a residential stressor with substantive consequences, greater disturbance should be associated with lower life satisfaction, more sleep difficulty, more depressive symptoms, and greater related health problems. Those outcomes with shorter-run response and recovery should also move in the adverse direction within the same respondent between waves of lower and higher reported disturbance.

Second, the timing of selection and reporting matters. If deteriorating well-being itself generates a new noise report, adverse outcome levels or recent adverse changes should predict the onset of high reported disturbance, and future disturbance may appear related to the current outcome even after current disturbance is taken into account. If residential selection is important, prior housing and neighbourhood circumstances should predict noise disturbance onset. Note that *contemporaneous* residential satisfaction indicators may deteriorate because the noise disturbance has begun so are generally not clean confounders.

Third, the framework predicts considerable heterogeneity grounded in the particular social circumstances; recurrent and night-time sounds, high transmission and limited control should generally carry a greater burden, while benign relationships, effective action or clearly bounded activity may reduce it.

The two components of our empirical research strategy address different sets of these predictions. The ELSOC panel estimates between- and within-person relationships, examines observed selection into reported disturbance and tests whether prior outcomes predict onset or continuation. Its annual measure cannot, however, identify the precise timing of individual acoustic events or separate sound pressure from appraisal. The Temuco interviews, on the other hand, do examine the nuanced social context: source, timing, recurrence, structural transmission, social meaning, perceived control, coping, household variation and the experience of obtaining or failing to obtain relief. Benign, mixed and discordant cases are retained because they identify the boundaries of the survey construct. While neither the quantitative nor the qualitative analysis by itself can entirely rule out confounding explanations, agreement across the two sources narrows the set of plausible interpretations.

3 Quantitative evidence

3.1 Data and analytic sample

The quantitative analysis uses the Chilean Longitudinal Social Survey (Estudio Longitudinal Social de Chile, ELSOC), conducted by the Centre for Social Conflict and Cohesion Studies (COES). ELSOC is a panel survey representative of Chile's national urban population. Its target population consists of men and women aged 18–75 who are usual residents of occupied private dwellings in urban areas; the design is probabilistic, stratified and multistage, with clustering at the primary-sampling-unit level and refreshment of the panel (Centro de Estudios de Conflicto y Cohesión Social (COES) 2025). We use seven survey waves: 2016, 2017, 2018, 2019, 2021, 2022 and 2023. Fieldwork paused in 2020 during the COVID-19 pandemic. Interviews were primarily face-to-face before the pause, by telephone in 2021, and mixed mode in 2022. Survey mode, sample membership, interview timing, original design strata and clusters, and the available cross-sectional and longitudinal weights are retained for diagnostic and sensitivity analyses.

The analysis-ready master file contains 20,761 person-wave observations from 4,447 respondents. It remains in long form, with one record per respondent in each observed wave. Negative non-response codes and values outside the documented substantive range are coded missing; an item not fielded in a wave is never recoded as zero. The neighbour-noise item has uneven wave-specific coverage, particularly in 2021 and 2022, so all models use outcome-specific available-case samples and display the resulting numbers of observations and respondents. Within-person models require at least two valid joint observations of noise and the relevant outcome. We do not impute missing outcomes or exposures. For the limited housing-history controls, the last valid ownership or neighbourhood-tenure report is carried forward when the contemporaneous item is absent; all other exclusions are specification-specific complete-case exclusions.

3.2 Exposure and outcomes

The primary exposure variable of interest records how frequently, during the previous 12 months, the respondent reports that they or another household member were annoyed or disturbed by noise arising as a problem with neighbours (with loud music and shouting given as examples) on a five-category scale from "never" to "always". It therefore measures experienced disturbance from neighbour noise rather than physical sound pressure. The ordinal measure is used for the main between–within analysis, so its coefficient is the change in the outcome associated with a one-category increase in reported noise. Non-linear fixed-effects models replace the linear term with category indicators. We also define high noise as reporting disturbance at least "sometimes" (categories 3–5) and low noise as "never" or "rarely" (categories 1–2) for the group profiles, binary fixed-effects comparison and exposure-prediction analyses. The alternative threshold of "many times" or "always" is used as a robustness check. As Section 2.4 distinguishes, the ELSOC exposure measures reported experienced disturbance rather than acoustic exposure. It captures a proximal residential stressor shaped by transmission, timing, recurrence and appraisal. The interviews later provide a more vivid sample of the range of sources, timing, appraisals and residential constraints represented by this scaled response.

The three primary outcomes are life satisfaction, sleep problems and depressive symptoms. Life satisfaction is measured by asking respondents: All things considered, how satisfied or dissatisfied are you with your life at the moment?", using a five-point scale from completely dissatisfied to completely satisfied. Sleep problems are measured using an item from the depressive-symptom battery referring to the previous two weeks. Respondents are asked how often they have experienced difficulty falling asleep, sleeping through the night, or sleeping too much", with responses coded from 1, never", through some days", more than half the days" and most days", to 5, "every day". Higher values therefore indicate more frequent sleep difficulty or excessive sleep. Depressive symptoms are measured using the remaining eight items in the same battery, covering loss of interest or pleasure, low mood or hopelessness, fatigue, appetite disturbance, difficulty concentrating, negative self-evaluation, observable psychomotor slowing, and thoughts of death or self-harm. We exclude the sleep item because it is analysed separately as a primary outcome and require valid responses to at least six of the eight retained items. The resulting measure is a symptom-frequency index rather than a clinical diagnosis. Secondary repeated outcomes include closeness of current life to the respondent's ideal, self-rated health, plans to move, and wanting to move while reporting being unable to do so.

3.3 Empirical strategy

We begin with descriptive comparisons among respondents who report low noise throughout 2016–2019, high noise throughout the same period, and those whose status changes. These groups show how persistent exposure and the within-estimation sample differ in observed composition, later retention and subsequent exposure. They are descriptive exposure histories rather than treatment groups.

The principal longitudinal decomposition is a within–between, or Mundlak, specification:

$$Y_{it} = \alpha + \beta_W (N_{it} - \bar{N}_i) + \beta_B \bar{N}_i + \lambda_t + \boldsymbol{\gamma}' \mathbf{X}_{it} + \varepsilon_{it}, \tag{1}$$

where Y_{it} is the outcome in its natural units for respondent i in survey wave t, N_{it} is reported neighbour noise, $\bar{N}_i$ is that respondent's mean reported noise over the observations used for the outcome, and λ_t denotes survey-wave effects. In the adjusted columns, $\mathbf{X}_{it}$ contains both respondent means and within-respondent deviations for every time-varying control, so that the between and within coefficients remain comparably adjusted. The between coefficient β_B compares people with different average reported exposure. It may combine long-running exposure with stable differences in housing, neighbourhood, vulnerability and reporting. The within coefficient β_W compares the same person in waves with more and less reported noise. It removes stable person-level selection but remains vulnerable to time-varying confounding, common-method reporting and measurement error.

The principal regression table reports each outcome in its original units and retains its natural direction: higher life satisfaction is better, while higher sleep and depressive-symptom scores are worse. For cross-outcome visual comparison and the robustness specifications, repeated outcomes are also standardised within their concurrent noise–outcome samples. Figure 1 sign-reverses life satisfaction, ideal life and self-rated health so that every plotted coefficient points in the adverse direction.

For the within-person estimates, we begin with a restricted specification and cumulatively add time-varying controls. The within-person component of the within–between model, and the equivalent respondent fixed-effects checks, remove all time-invariant observed and unobserved characteristics from the comparison, including sex, stable educational attainment and stable response dispositions. Survey-wave effects absorb shocks common to respondents in the same ELSOC wave, including common economic cycles. The first column therefore includes respondent and survey-wave effects but no additional time-varying controls. The second column adds within-person changes in log nominal household income and the respondent's position on the consistently fielded subjective-status ladder. Education is not added as a time-varying control: it is fixed for many adults, while its remaining within-person variation principally records educational progression and requires a separate life-course interpretation.

The third column adds within-person variation in the most recent reported home-ownership status and years in the neighbourhood. Ownership and neighbourhood tenure are not asked in every wave, so the last past report is carried forward when the contemporaneous module is absent; future values are never used. Because neighbourhood tenure is observed only in 2016, 2018 and 2023, respondent and wave effects absorb the common increase in reported tenure between fielding waves. The remaining variation reflects respondent-specific revisions in reported tenure, potentially arising from residential moves, changes in the geographical area respondents regard as their neighbourhood, rounding, or reporting error. We therefore treat

tenure as an expansive residential-history robustness control rather than a direct measure of moving. Contemporaneous home satisfaction, neighbourhood safety and trust, moving plans, and the desire to move while unable are not included in the principal control blocks. They may reflect a common change in residential circumstances, but they may also be consequences or mediators of neighbour noise. We instead examine them as outcomes, predictors measured prior to a transition, or deliberately severe sensitivity controls.

The between-person coefficient compares respondents with different average levels of reported neighbour noise over the same outcome-specific estimation sample. The first column is an unadjusted benchmark apart from survey-wave effects. The second column adjusts the between comparison for sex and educational attainment at the first valid report, age normalised to 2016, and the respondent-level means of log household income and subjective status. The third column additionally controls for the respondent-level means of the carried-forward ownership and neighbourhood-tenure measures. Thus the progressively richer columns use the same broad socioeconomic and residential domains for both components, while respecting their different sources of identifying variation: changes around an individual's own mean for the within estimate, and differences in average characteristics across individuals for the between estimate. The main estimates are followed by tests of whether prior outcomes predict onset or persistence and whether future noise predicts current outcomes, omitted-variable sensitivity calculations, heterogeneity checks and a separate analysis of incident diagnoses. The relevant samples and specifications are described with their results. Standard errors are clustered by respondent throughout. We report unadjusted p-values and apply the Holm procedure within prespecified primary- and secondary-outcome families for each principal estimand.

For robustness we compare the within component with conventional respondent fixed-effects estimates and adjacent-wave first differences. First-difference models condition on the starting noise category and survey wave; a separate specification excludes the two-year 2019–2021 interval. Non-linear category models, alternative binary thresholds, stable-commune samples, panel-length restrictions, fieldwork-mode checks and influence diagnostics probe whether the result is an artefact of the linear scale, a particular wave or a small group of high-leverage respondents. These checks are collected in Appendix Table 7.

All analyses were conducted in Stata 18; full replication materials are available at https://github.com/dmweinhold/Neighbour_Noise_Chile_Replication.

3.4 Neighbour noise in the panel

The first point to emerge from the data is that neighbour noise is not well described as a fixed attribute of the respondent, nor is the observed variation confined to a small set of households. Among the 4,020 respondents for whom we observe neighbour noise in at least two waves, 61.7 percent cross the binary high-noise threshold at least once and 85.9 percent change their position on the original ordinal scale. The probability of ever changing is, of course, partly mechanical: respondents observed in seven waves have many more opportunities to change than respondents observed twice. Nevertheless, the underlying within-person variance is widely distributed. The most variable 10 percent of respondents account for 30.7 percent of the within-person sum of squares. Thus the fixed-effects results below are not generated by a very small and obviously exceptional set of households.

For some of the analysis it is useful to characterise respondents by their early noise history. Requiring valid observations throughout 2016–2019 produces 687 respondents who remain at low or no noise, 116 who remain at high noise, and 1,108 whose status changes. The persistently

exposed group continues to report appreciably more noise in the later waves: 53.8 percent report high noise in 2023, compared with 20.6 percent of the persistently low group. The three groups have very similar subsequent survey retention, at approximately 77 percent. This is reassuring, since differential exit from the panel is one of the more immediate ways in which a comparison of persistent exposure histories could become misleading.

Table 1: Characteristics in 2016 by 2016–2019 neighbour-noise history

Characteristic	Persistent low	Changer	Persistent high	Std. high–low	p-value
Panel A. Demographic characteristics and conventional socioeconomic position					
Age (years)	48.78	46.43	45.15	−0.24	0.013
Male (percent)	39.2	36.7	32.8	−0.13	0.178
Education level (1–10)	5.21	5.12	5.27	0.03	0.791
Completed university (percent)	11.8	10.7	10.3	−0.05	0.640
Log household income	12.96	12.96	13.01	0.06	0.527
Subjective status (0–10)	4.43	4.24	4.10	−0.21	0.020
Subjective class (1–5)	2.31	2.24	2.17	−0.17	0.072
Unemployed (percent)	5.5	5.0	7.9	0.11	0.377
No savings (percent)	55.9	57.0	59.1	0.06	0.529
Panel B. Housing position and residential history					
Home owner (percent)	72.3	69.3	65.5	−0.15	0.152
Years in neighbourhood	20.99	19.82	17.07	−0.24	0.007
Panel C. Contemporaneous residential appraisal and mobility					
Home-quality satisfaction (1–5)	3.76	3.61	3.48	−0.29	0.009
Neighbourhood safety (1–5)	3.65	3.31	2.65	−0.96	< 0.001
Trust in neighbours (1–5)	3.56	3.09	2.56	−0.84	< 0.001
Plans to move (percent)	14.1	20.1	32.8	0.48	< 0.001
Wants to move but cannot (percent)	15.3	23.6	35.3	0.49	< 0.001

Notes: Entries are unweighted 2016 means. Persistent-low respondents report low noise ("never" or "rarely") in every wave from 2016 to 2019; persistent-high respondents report high noise ("sometimes" or more) throughout; changers cross this threshold at least once. Group sizes are 687 persistent-low, 1,108 changer and 116 persistent-high respondents. Item-specific sample sizes range from 552 to 687 for persistent low, 896 to 1,108 for changers, and 95 to 116 for persistent high. The standardised difference and unadjusted two-sided p-value compare persistent high with persistent low. Because 2016 is the first year of the exposure-classification window, these are contemporaneous profiles, not genuinely pre-exposure predictors. Panel C variables may be antecedents, consequences or mediators of neighbour noise and are therefore not included in the preferred control block.

Table 1 explores possible selection profiles. Education, university completion, household income, unemployment and financial precarity are nearly identical across the three noise exposure groups. Persistent-high respondents are somewhat younger and report modestly lower subjective status, but there is no strong conventional socioeconomic gradient. Home ownership is also similar, although persistent-high respondents have spent fewer years in the neighbourhood. By contrast, their reported neighbourhood safety is 0.96 standard deviations lower than that of the persistent-low group, neighbour trust is 0.84 standard deviations lower, and plans or constrained desires to move are approximately 19–20 percentage points more common. Changers are generally intermediate.

These patterns make a simple low-income or low-education sorting explanation difficult to sustain. The 2016 profiles are contemporaneous with the first year used to define the groups, so safety, trust, home satisfaction and moving intentions may reflect antecedent residential conditions, responses to noise, or both. A problem with neighbour noise exposure may precipitate or heighten all these negative assessments, so predetermined housing history, including contemporaneous residential appraisal or a wish to move, are potentially at least partially confounding. The qualitative interviews underscore this caution: dissatisfaction and moving intentions are frequently described as responses to noise disturbance (Section 4).

3.5 Between- and within-person relationships

Table 2 reports the principal estimates for life satisfaction, sleep problems and depressive symptoms, with the depressive-symptom measure constructed to exclude the sleep item. The outcomes retain their original units and natural signs. Panel A compares people with different average noise exposure, while Panel B compares the same respondent in waves in which he or she reports more rather than less neighbour noise. All three columns use the same outcome-specific sample. For the within-person component, the second column adds household income and subjective status and the third adds ownership and neighbourhood-tenure history. For the between-person component, the second column additionally adjusts for sex, age and education. Contemporaneous subjective assessments of the residential environment are excluded throughout.

Table 2: Between- and within-person associations under progressively richer control sets

	Wave effects only	+ Demographic and socioeconomic controls	+ Residential history
Panel A. Between-person coefficient			
Life satisfaction	−0.106	−0.086	−0.085
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Sleep problems	0.186	0.169	0.168
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Depressive symptoms	0.112	0.099	0.099
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Panel B. Within-person coefficient			
Life satisfaction	−0.0248	−0.0237	−0.0237
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Sleep problems	0.0516	0.0513	0.0510
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Depressive symptoms	0.0376	0.0371	0.0369
	($p < 0.001$)	($p < 0.001$)	($p < 0.001$)
Observations	16,874–16,882	16,874–16,882	16,874–16,882
Respondents	4,267–4,268	4,267–4,268	4,267–4,268

Notes: Entries are changes in the outcome's original units per one-category increase in reported neighbour noise; unadjusted two-sided p-values are shown beneath the estimates. Life satisfaction and sleep problems are measured from 1 to 5. Depressive symptoms are the mean of eight 1–5 items, excluding sleep. Thus a negative life-satisfaction coefficient and positive sleep or depression coefficients indicate adverse outcomes. All columns use the same outcome-specific complete-case sample and include survey-wave effects. The correlated-random-effects specification decomposes noise and every time-varying control into respondent means and within-respondent deviations. Column 2 adds log household income and subjective status to both components and additionally adjusts the between-person component for sex, age normalised to 2016, and education at the first valid report. Column 3 adds the most recent reported home-ownership status and years in the neighbourhood, decomposed into their between- and within-person components. Standard errors are clustered by respondent.

The between-person estimates are large and remarkably consistent across outcomes. In the first column, a one-category increase in average neighbour noise is associated with a 0.106-point reduction in life satisfaction, a 0.186-point increase in sleep problems and a 0.112-point increase in mean depressive symptoms. Adding the full demographic, socioeconomic and residential-history control block only slightly attenuates these coefficients to −0.085, 0.168 and 0.099. Most of the attenuation occurs when age, sex, education, household income and subjective status are introduced; ownership and neighbourhood-tenure history add little thereafter.

The within-person results point in the same direction. When an individual reports a one-

category increase in neighbour noise, life satisfaction falls by 0.0248 points, sleep difficulty rises by 0.0516 points and mean depressive symptoms rise by 0.0376 points. The fully adjusted estimates are nearly identical, at −0.0237, 0.0510 and 0.0369, respectively, and all have $p < 0.001$. Expressed in standard-deviation units, the fully adjusted headline specification associates moving from low to high reported noise with a 0.064 SD fall in life satisfaction, a 0.065 SD increase in sleep problems and a 0.102 SD increase in depressive symptoms. To put these changes in perspective, the high-noise association is approximately 54 percent as large as becoming unemployed for life satisfaction, similar to becoming unemployed for sleep problems, and slightly larger than becoming unemployed for depressive symptoms. Becoming romantically partnered is associated with a 0.082 SD improvement in life satisfaction, while transition from low to high noise is associated with a similar magnitude (0.078) SD reduction.

Observed changes in socioeconomic position and residential history account for little of the within-person relationships. Appendix Table 7 shows the same conclusion under conventional respondent fixed effects and first differences, exclusion of the telephone wave, restriction to stable observed communes and longer panels, removal of high-leverage exposure histories, alternative binary thresholds and a fully non-linear exposure specification. The coefficient signs are unchanged throughout, and each primary within-person relationship remains statistically significant.

Figure 1 places the primary and secondary outcomes on the same adverse-direction scale (so an increase indicates a favorable change). The much larger between-person coefficients are immediately visible, but so is the consistency of the within-person direction. Plans to move and wanting to move but being unable to do so are among the strongest secondary associations. As discussed, the qualitative cases suggest why these measures are especially difficult to regard as clean controls: some households begin considering a move only after repeated noise has made the dwelling difficult to inhabit.

To gauge sensitivity to omitted variables, we apply the Diegert–Masten–Poirier framework to outcomes, exposure and the full nested control block residualised with respect to respondent and wave effects. The resulting breakdown frontiers ask how strongly a time-varying omitted factor would have to relate to both reported noise and the outcome, relative to observed changes in income, subjective status and residential history, to permit a sign reversal. Appendix Table 7, Panel C, reports selected points along the frontier. At the benchmark maximum correlation of 0.5 between the omitted factor and the observed controls, and assuming equal relative strength in the exposure and outcome equations, a time-varying confounder would have to be 1.33, 1.62 and 1.57 times as influential as the observed income, subjective-status and residential-history block to permit a sign reversal in the life-satisfaction, sleep and depression estimates.

3.6 Selection and the direction of timing

The data allows us to examine the leading alternative explanations for the noise correlations reported above. We model onset among previously low-noise respondents and persistence among previously high-noise respondents as functions of prior outcomes, prior outcome changes, socioeconomic circumstances, residential context and other reported neighbourhood problems. Future-noise placebo models ask whether next-wave noise predicts the current outcome once current noise and person and wave fixed effects are included. Prespecified interaction models consider age, sex, baseline socioeconomic position, home ownership, nervousness and reported calmness under stress; person-level response plots assess whether within-person relationships separate into visibly distinct clusters.

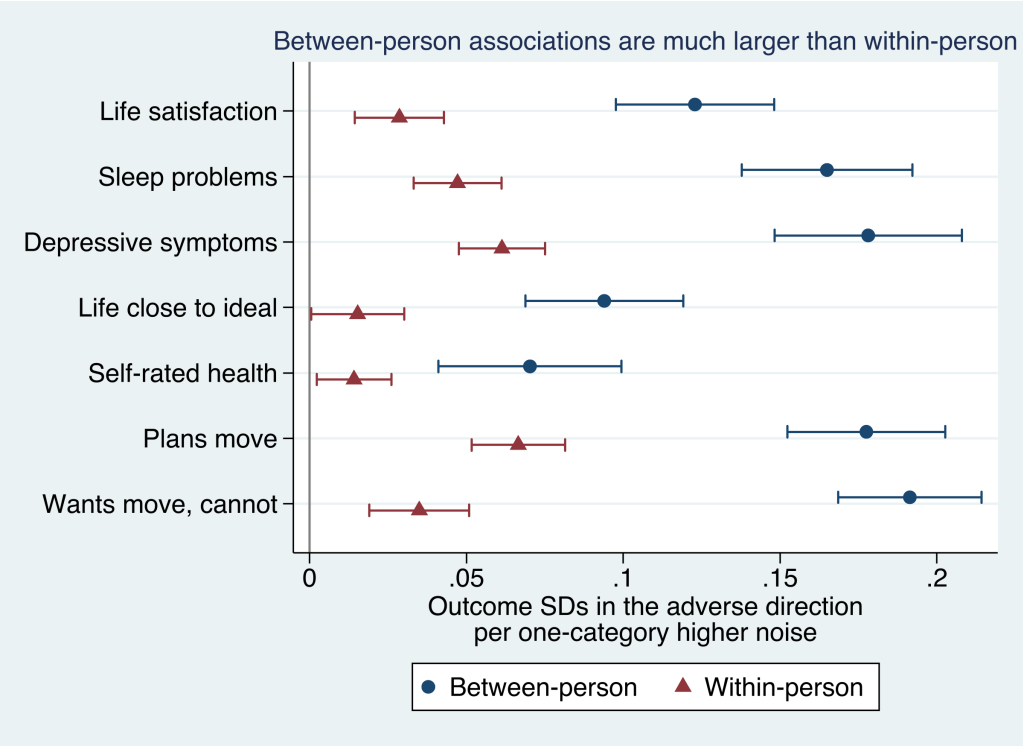

Figure 1: Between- and within-person associations between neighbour noise and well-being, health and mobility outcomes. Points are standard-deviation changes in the adverse direction per one-category increase in neighbour noise; bars show 95-percent confidence intervals. Blue circles are the between-person component and red triangles the within-person component from the unadjusted correlated-random-effects model on each outcome's full available sample. Life satisfaction, life close to ideal and self-rated health are sign-reversed. Table 2 reports the primary outcomes in their original units on the common nested-control sample.

If respondents who are becoming depressed or losing sleep subsequently report more noise, the within-person association could still run in the opposite direction. Table 3 collects the tests of this interpretation and of selection into exposure. Indeed, onset is not random: a one-standard-deviation deterioration in prior residential context predicts 3.40 percentage points more onset, while a one-standard-deviation increase in other reported neighbourhood problems predicts 4.07 percentage points more onset (Panel A). Nevertheless the temporal ordering is reassuring for the primary outcomes. None of their prior adverse levels significantly predicts onset after adjustment across the three tests (Panel B). Recent adverse changes also do not predict onset in the expected direction: all three coefficients are negative, and the statistically significant depressive-symptom coefficient therefore runs against a reverse-reporting account. In the future-noise placebo, the coefficients on next-wave noise are essentially zero once current noise, respondent fixed effects and survey-wave effects are included (Panel D). These results weaken a simple explanation in which worsening well-being causes respondents newly to report noise.

There is one important qualification. Among respondents already reporting high noise, one standard deviation worse prior sleep predicts a 4.67-percentage-point greater probability of another high report (Holm-adjusted $p < 0.001$; Panel C). This could reflect reporting heterogeneity, a persistently more disruptive exposure, or both, and is a reason to interpret chronic exposure more cautiously than onset. The negative coefficients on recent outcome changes in this panel indicate that short-run deterioration itself is not what predicts continued reporting.

Table 3: Selection into high reported noise and reverse-timing tests

Panel A. Prior residential and mobility characteristics predicting next-wave onset			
Prior characteristic	Effect on onset (percentage points)	Person-waves	Respondents
Residential context (1 SD better)	−3.40 $p < 0.001$	8,285	3,257
Other neighbourhood problems (1 SD worse)	4.07 $p < 0.001$	8,285	3,257
Plans to move	0.74 $p = 0.587$	8,285	3,257
Wants to move but cannot	3.18 $p = 0.012$	8,285	3,257

Panel B. Prior outcomes predicting next-wave onset				
Outcome	Prior adverse level	Prior adverse change	Person-waves	Respondents
Life satisfaction	0.26 $p = 0.735$	−1.21 $p = 0.092$	5,552	2,355
Sleep problems	1.34 $p = 0.137$	−0.92 $p = 0.133$	5,553	2,354
Depressive symptoms	1.47 $p = 0.128$	−1.85 $p = 0.015$	5,558	2,354

Panel C. Prior outcomes predicting continued high reporting				
Outcome	Prior adverse level	Prior adverse change	Person-waves	Respondents
Life satisfaction	2.43 $p = 0.086$	−0.93 $p = 0.353$	2,465	1,436
Sleep problems	4.67 $p < 0.001$	−2.83 $p = 0.019$	2,467	1,440
Depressive symptoms	2.37 $p = 0.086$	−2.13 $p = 0.098$	2,463	1,438

Panel D. Future-noise placebo: next-wave noise predicting the current outcome			
Outcome	Coefficient on next-wave noise	Person-waves	Respondents
Life satisfaction	−0.0079 $p = 1.000$	13,099	3,894
Sleep problems	−0.0011 $p = 1.000$	13,098	3,895
Depressive symptoms	−0.0080 $p = 1.000$	13,093	3,890

Notes: Panel A reports linear-probability estimates among respondents whose preceding report was low noise. The residential-context index combines home satisfaction, neighbourhood safety and neighbour trust; a larger value denotes a better context. The other-problems index combines reported problems with animals, threats and litter. The model also conditions on the exact prior noise category, age, sex, socioeconomic status and survey wave. Panels B and C report percentage-point changes in the probability of onset or continued high reporting for a one-standard-deviation worse prior outcome level or prior outcome change. These models additionally condition on the residential and mobility variables shown in Panel A; their p-values are Holm-adjusted across the three outcomes within each test. Panel D reports outcome-standard-deviation coefficients per one-category increase in next-wave noise, conditional on current noise, respondent fixed effects and survey-wave effects; its p-values are likewise Holm-adjusted. Standard errors are clustered by respondent throughout. High noise is defined as reporting noisy neighbours sometimes, frequently or always.

Prior plans to move do not predict noise onset, whereas previously wanting to move but being unable to do so predicts 3.18 percentage points more onset. At the same time, current reported noise is strongly associated with both outcomes. They plausibly capture residential constraint and selection, but they can also arise as behavioural responses to noise and to feeling trapped by it. We therefore report the lagged association but do not treat these variables as clean contemporaneous controls. The qualitative evidence makes the latter ordering explicit in cases where repeated disturbance precedes discussion of leaving the dwelling.

Relocation could also pose a selection issue. Working with Dutch survey data, Cassens et al. (2025) note that their models cannot control for moves and take the view that relocation is unlikely to occur systematically. Yet moving can follow neighbour noise, and the events that precipitate a move can independently affect well-being. Based on an earlier sample of the same

Dutch data but augmented with moving data, Weinhold (2015) finds that neither current nor past neighbour noise predicts moving and that controlling for moves leaves the health coefficients unchanged. In the Chilean ELSOC data, restricting the sample to respondents whose observed commune remains stable produces within-person estimates very close to those in the full sample (Appendix Table 7). This is a blunt check because it captures cross-commune changes rather than local moves, but it gives no indication that observed long-distance mobility drives the principal relationships.

Observable subgroup interactions provide little evidence of stable "perturbable" and "imperturbable" types, and individual-response plots are dispersed rather than cleanly clustered. Exploratory decomposition of baseline socioeconomic position produces a suggestive but non-uniform pattern. Income and education are associated with weaker sleep and depressive-symptom responses to noise, whereas subjective class and subjective status more strongly moderate life satisfaction. Only the interaction between subjective class and life satisfaction survive adjustment for multiple testing. Direct comparisons between the material and subjective interaction coefficients are nevertheless imprecise, and analogous checks using closeness of life to the ideal and self-rated health does not reproduce a clean division between evaluative and non-evaluative outcomes. The surviving life-satisfaction result may therefore reflect differences in appraisal or scale use as well as substantive differences in resilience. We treat these decompositions as hypothesis-generating rather than as evidence of an established socioeconomic gradient, and report the full specification family in Appendix Table 8.

Baseline income and education nevertheless differ little across persistent-low, changing and persistent-high exposure histories, so what material inequality there may be seems more likely to lie in households' capacity to respond to noise than in exposure to it. This is a claim about conventional material position rather than about socioeconomic inequality in general: as Table 1 shows, persistently exposed respondents do report lower subjective status, shorter neighbourhood tenure and substantially worse residential appraisals. Again, the interviews sharpen the interpretation: sensitivity is often relational and situational, and members of the same household can respond quite differently to the same source.

3.7 Incident physical-health diagnoses

Theory and some suggestive evidence discussed above suggests that neighbour noise could be associated with a number of diseases, moderated via both sleep disruption and/or stress. In the Chilean data chronic diagnoses are analysed separately because the diagnosis battery was administered only in 2017 and 2022. Respondents reporting no baseline diagnosis in 2017 enter diagnosis-specific risk sets, and the outcome is a newly reported diagnosis dated 2020–2022. Exposure is measured before that window, over 2016–2019. The primary minimum baseline ages are 50 for heart disease, 40 for hypertension, 35 for diabetes, 18 for diagnosed depression and 50 for cancer; neighbouring thresholds are reported as sensitivity checks rather than searched for favourable results. Linear-probability models adjust for baseline age, sex and socioeconomic position, with diagnosis-specific baseline risks added only in sensitivity models. Bias-reduced Firth logistic models and exact comparisons are used where event counts are sparse. For the persistent-group comparisons, observable standardisation compares persistent-high and persistent-low respondents over common covariate support.

Figure 2 displays the persistent-group comparison across all five possible diagnoses. These analyses are explicitly exploratory; new diagnoses are rare, their plausible risk populations differ, and the survey waves provide only a limited window in which both exposure and incidence can be observed. Other than for heart disease the group comparisons are not statistically significant,

although depression displays the expected monotone increase. The heart result is is sparse, but striking. Incident heart disease occurs in 5 of 238 persistently low respondents (2.1 percent), 12 of 373 changers (3.2 percent), and 4 of 34 persistently high respondents (11.8 percent). The adjusted linear-probability difference between the high- and low-noise groups is 9.93 percentage points (95-percent CI −0.95 to 20.81), while the Firth odds ratio is 7.29 (95-percent CI 1.89 to 28.08). Each of the four high-noise cases therefore matters. Reweighting to common covariate support reduces the high-versus-low estimate and makes it imprecise, partly because the effective number of persistently exposed respondents falls to approximately 16. Removing any one of the four cases also weakens the result. Thus the appropriate conclusion is that the continuous and grouped analyses reveal the same unusual concentration of incident heart diagnoses among older respondents with high exposure, but on a very limited sample.

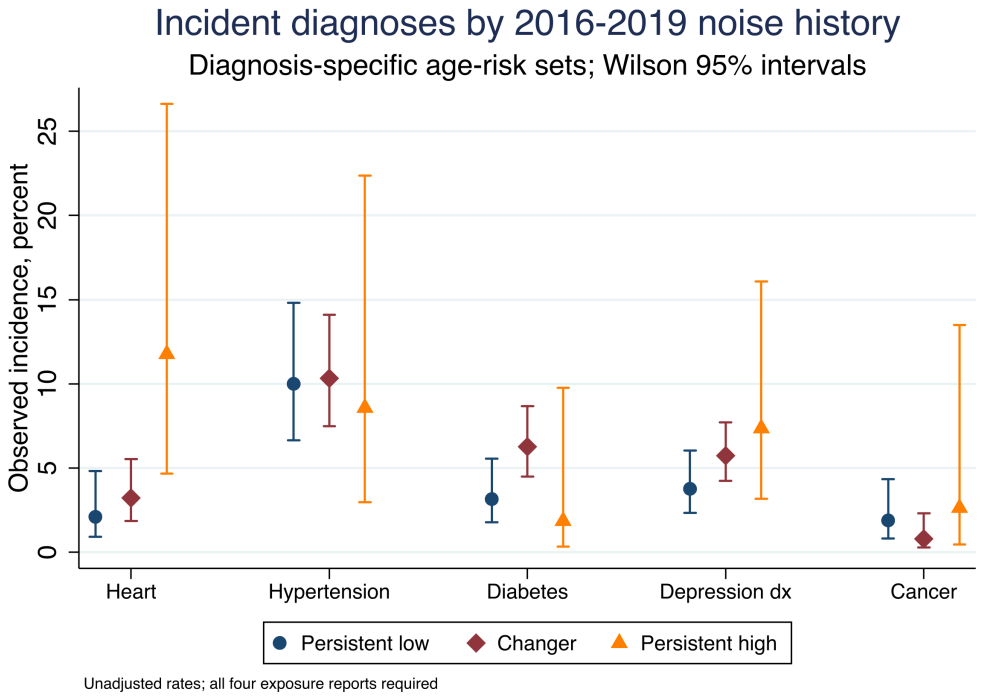

Figure 2: Unadjusted incident-diagnosis rates by 2016–2019 neighbour-noise history. Bars show Wilson 95-percent intervals in diagnosis-specific, age-appropriate risk sets. All four early exposure reports are required. The heart-disease contrast is visually large but is based on only four cases among 34 persistently high respondents; adjusted heart and hypertension contrasts are reported in Table 4.

Table 4 reports a more detailed exploration of the heart-disease estimates; we present these alongside hypertension, which provides both a wider cardiovascular risk set and a useful null comparison.

For heart disease, the principal age-50-and-over risk set contains 686 respondents and 23 incident diagnoses during 2020–2022. A one-category increase in mean noise during 2016–2019 is associated with a 1.72-percentage-point increase in subsequent diagnosis (95-percent CI −0.01 to 3.46; $p = 0.052$). When early and recent exposure are entered together, the association is concentrated in 2018–2019: a one-category increase in recent noise predicts 2.18 percentage points more incident heart disease (95-percent CI 0.22 to 4.14; $p = 0.030$), while the earlier exposure coefficient is close to zero. The rare-event Firth model gives a similar result (OR 1.74, 95-percent CI 1.13 to 2.66). The estimate is also similar when the age threshold is varied from 45 to 55. It does not, however, survive correction across every diagnosis, exposure window and

Table 4: Exploratory incident-diagnosis results in diagnosis-appropriate risk sets

Diagnosis	Exposure or contrast	Estimate	95% CI	p-value
Panel A: Continuous early exposure				
Heart disease	Mean noise, 2016–2019 (pp)	1.72	[−0.01, 3.46]	0.052
Heart disease	Recent noise, 2018–2019 (pp)	2.18	[0.22, 4.14]	0.030
Heart disease	Recent noise, Firth odds ratio	1.74	[1.13, 2.66]	0.012
Hypertension	Mean noise, 2016–2019 (pp)	0.43	[−2.26, 3.12]	0.754
Hypertension	Recent noise, 2018–2019 (pp)	−1.35	[−4.28, 1.58]	0.366
Hypertension	Recent noise, Firth odds ratio	0.87	[0.63, 1.19]	0.373
Panel B: Persistent early-exposure histories				
Heart disease	Persistent low: 5/238	2.1%		
	Changer: 12/373	3.2%		
	Persistent high: 4/34	11.8%		
	Adjusted high minus low (pp)	9.93	[−0.95, 20.81]	0.074
	Firth high versus low odds ratio	7.29	[1.89, 28.08]	0.004
Hypertension	Persistent low: 21/210	10.0%		
	Changer: 34/329	10.3%		
	Persistent high: 3/35	8.6%		
	Adjusted high minus low (pp)	−1.39	[−11.52, 8.74]	0.787

Notes: "pp" denotes percentage points per one-category increase in neighbour noise. The primary risk set is age 50 or over for heart disease and age 40 or over for hypertension, with respondents diagnosed by 2017 excluded. Panel A contains 23 heart-disease events in 686 respondents for mean exposure (21/641 for the recent-exposure model) and 61 hypertension events in 608 respondents (58/570 for recent exposure). Panel B uses respondents observed in every 2016–2019 exposure wave; adjusted contrasts use the compact baseline risk and socioeconomic controls. The Firth odds ratios are shown because the persistent-high heart group contains only four cases. These estimates are exploratory and do not survive correction across the full diagnosis-by-window specification family.

model examined, and should not be read as a confirmatory test.

This pattern is notably similar to that in the earlier Netherlands study, despite the different country, period and survey (Weinhold 2015). We also find limited evidence that sleep is a pathway: adding baseline and intermediate sleep measures attenuates the mean-noise coefficient by about 12 percent, but the estimated attenuation is itself imprecise. Stress remains an equally plausible mechanism in these data. By contrast, high blood pressure—for which the age-40-and-over risk set includes 61 incident diagnoses—has no stable relationship with noise in either the continuous or persistent-group models. The more weakly suggestive diabetes and diagnosed-depression results are best treated as exploratory, while the cancer sample is too small to interpret.

The qualitative interviews do not collect health information and clinically interpretable physical-health statements are rare, so they are primarily informative here for better understanding potential mechanisms. Sleep disturbance, next-day fatigue, hypervigilance and anticipatory stress are common among the respondents, supporting the plausibility of sleep and stress pathways. Taken together, the heart-disease result is interesting but remains a hypothesis-generating result rather than a finding due to the limited amount of data.

4 Qualitative evidence

4.1 Setting, sampling and recruitment

The qualitative study was conducted in Temuco, the capital of the Araucanía Region in southern Chile. Temuco provides substantial variation in dwelling form, tenure, density and neighbourhood context within a single institutional setting. The study was not designed to estimate prevalence, but to examine how ordinary domestic sounds cross dwelling boundaries and become interpreted as tolerable activity, nuisance, disrespect or threat. Thus we purposively recruited adults aged 18 or older who had lived in Temuco for at least twelve months and could discuss current or past experience of noise in the home. Recruitment was not random; invitations were circulated through online social networks and direct contacts, with snowball referrals used to extend coverage. We sought variation across city areas, dwelling types, tenure, household composition and basic housing characteristics. Because network-based and topic-based recruitment can over-represent highly sensitised respondents, recruitment and analysis actively sought contrasting and benign accounts as well as serious disturbance. Interviews continued until the team judged that additional cases were no longer adding materially different mechanisms or contextual conditions, rather than until a numerical frequency target was reached (Hennink et al. 2017).

The final analytic corpus contains 45 independent interview or household sessions. Two were multi-resident sessions: AJ5 included two participants and AJ45 three. Thus at least 48 people are explicitly recorded, but the pre-interview form and coded material supply one analytic cluster per session. Each multi-resident interview is consequently counted once, and sensitivity analyses exclude both[1].

4.2 Data collection

Fieldwork took place between December 2024 and April 2025, spanning ordinary weeks and the summer holiday period. Interviews followed a semi-structured guide covering the source, acoustic character, timing and recurrence of noise; emotional and bodily responses; perceived controllability, intentionality and relations with neighbours; sleep and next-day functioning; coping and adaptation; and experiences of direct or institutional complaint. Participants were encouraged to recount concrete episodes, distinguish current from past dwellings and compare daytime with night-time disturbance.

A short pre-interview questionnaire recorded demographic and housing information, including tenure, dwelling type, approximate construction year, glazing and household composition, together with closed-form ratings of nuisance from neighbour and other urban-noise sources. Interviews were conducted in Spanish, audio-recorded with consent, transcribed and de-identified. Potentially identifying details were removed from quotations. All participants provided informed consent. The protocol received approval from the London School of Economics and Political Science Research Ethics Committee on 25 October 2024.

[1] A further source identifier, AJ1, contains no transcript or substantive episode and is excluded from the analytic corpus. This reconciliation explains the difference between the number of source identifiers and the final number of analytic sessions.

4.3 Thematic and case-oriented analysis

The interview strand uses an abductive, reflexive thematic analysis (Braun & Clarke 2006, 2019). An initial codebook derived from the research questions, prior noise research and pilot readings was revised as new patterns, counter-examples and construct boundaries appeared. The coding combined a priori themes—including timing, recurrence, sleep, stress, perceived control, coping and institutional response—with inductive codes for unanticipated meanings and experiences. Transcript organisation and thematic coding used ATLAS.ti. Analysis proceeded through constant comparison within and across cases, matrix-based contrasts and examination of co-occurring themes rather than through frequency counts alone (Miles et al. 2014). Regular analytic discussions were used to challenge assumptions about reasonable quiet, neighbourliness and what constitutes successful adaptation.

This human-led interpretive analysis was supplemented by a structured coding workflow designed to make comparisons across interviews and with the ELSOC survey data more systematic. We describe the workflow as *LLM-assisted, human-adjudicated qualitative content analysis*. The approach was deliberately narrower than "automated thematic analysis". Recent work shows that large language models can help propose codebooks, recover concrete descriptive categories and apply pre-specified codes, but are less dependable when the task requires latent meaning, contextual judgement or theoretical synthesis (Morgan 2023, Hamilton et al. 2023, Dai et al. 2023, Wachinger et al. 2025). In our design the model therefore organised candidate evidence; it did not on its own determine the final themes or interpretation. A detailed account is provided in Appendix Appendix B. The remainder of this section reports what the combined case-oriented and structured analysis reveals, beginning with the corpus and the unit at which comparisons can be interpreted.

4.4 The corpus and unit of interpretation

The interviews provide a different kind of evidence from the panel. They cannot supply a counterfactual or a population prevalence estimate, but they can show what respondents mean when they report noise, how a disturbance enters daily life, and which parts of the quantitative interpretation are more or less plausible. We treat the 45 interview or household sessions described above as independent analytic clusters. The structured material contains 343 substantive episodes, but these episodes are nested mentions rather than independent observations; all descriptive comparisons below therefore return to the interview or household session as the unit.

Applying the neighbour-episode definition described above yields 106 episodes distributed across 43 of the 45 sessions. The high fraction is unsurprising given the recruitment and topic and must not be read as the prevalence of neighbour noise in Chile. Its value is instead that the corpus contains considerable variation in form, timing, appraisal and response among people who can all describe some relevant exposure.

Table 5 summarises the main coded pathways and outcomes. The coded episode dataset records pathway labels, which identify the mechanism linked to a trigger, separately from outcome labels, which identify the response or endpoint; thematic counts throughout take the union of the two where a construct appears in both. Themes are not mutually exclusive, and 23 of the 43 neighbour-episode sessions contain both an adverse response and a benign or no-problem appraisal somewhere in the interview. People do not divide neatly into those disturbed by all noise and those disturbed by none. In six sessions the two appraisals attach to the same sound, which is a stronger and rarer form of the same point.

Table 5: Qualitative themes among interview sessions containing strict-neighbour episodes

Theme or feature	Sessions	Percent	Quantitative complement
Panel A: Reported pathways and outcomes			
Annoyance or irritation	37	86.0	Contemporaneous well-being and life satisfaction
Benign appraisal or no problem	27	62.8	Resilience, adaptation and heterogeneous response
Stress or hypervigilance	25	58.1	Stress pathway; depressive symptoms and possible cardiovascular mechanism
Sleep disruption	24	55.8	Sleep outcome, next-day fatigue and irritability
Social conflict	20	46.5	Neighbour relations, trust and residential context
Fear or insecurity	11	25.6	Threat appraisal and constrained confrontation
Concentration or productivity	10	23.3	Functional impairment
Hopelessness or demoralisation	3	7.0	Depressive symptoms and lower well-being
Self-described health symptom	2	4.7	Mechanism evidence only; not a clinical diagnosis
Panel B: Form, timing and response			
Any recurring strict-neighbour pattern	39	90.7	Persistent exposure rather than an isolated incident
Any night-time strict-neighbour episode	35	81.4	Sleep interruption and heightened night-time salience
Any day-time strict-neighbour episode	30	69.8	Annoyance and concentration during ordinary activity
Any constraint	13	30.2	Fear, social norms or an ineffective institution
Any agency	19	44.2	Direct confrontation or an effective institution
Adaptation without changing the source	17	39.5	Habituation or resignation, not necessarily absence of burden

Notes: Denominator $N = 43$ interview/household sessions containing at least one strict-neighbour episode. Themes are not mutually exclusive; in particular, adverse and benign appraisals frequently occur in the same interview. Sessions, not the 106 strict-neighbour episodes, are the counting unit. The interviews were purposively recruited and these percentages are descriptive of this corpus, not population prevalence estimates.

4.5 Sources, timing and acoustic texture

Reports cluster around a limited set of sources. By frequency of mention across all 343 episodes, road traffic is the most common trigger category (62 episodes), followed by neighbour parties (41), dogs (36), construction and power tools (31), neighbour voices (19), televisions and audio equipment (19), music (15) and fireworks (15).

The timing of a sound matters as much as its source, and the two are not independent. Daytime accounts are dominated by routine and functional sources: machinery, street vendors, school playgrounds, ongoing construction. Night-time concentrates sounds that residents read as discretionary. Among the 85 neighbour episodes carrying a day or night code, 33 occur only at night, 26 only during the day and 26 at both times; a further 21 episodes are not classified by time.

The clearest illustration of how far appraisal departs from acoustics comes from a source category that divides internally. Road vehicles are the most frequently mentioned trigger of all, but they split sharply by time and content. Daytime traffic appears as *tráfico vehicular de la avenida*, *locomoción y buses*, and in one respondent's phrase *ruidos normales de autos*. Night-time traffic is something else: *carreras clandestinas*, *tubos de escape modificados*, *autos robados pasando rajando*. Classifying each session by all of its traffic episodes, eight report night-time vehicle noise only, eight daytime only, sixteen both and six neither. Among the mutually exclusive groups, four of the eight night-only sessions record fear or insecurity, against none of the eight day-only sessions (two-sided Fisher exact $p = 0.077$, and $p = 0.026$ when

episodes flagged in the coding audit are excluded). The contrast is suggestive rather than conclusive on these numbers, but it points in a consistent direction: the same nominal source, at plausibly higher sound pressure at night, produces different responses, and what separates them is not the decibel level but whether the sound is read as somebody's choice.

Respondents also make fine acoustic distinctions that map onto the transmission literature. Low-frequency rumble, *retumbe*, is described as invasive and hard to escape, and appears in six episodes across four sessions: a washing machine running upstairs late at night, a lift mechanism audible through the structure, modified speakers in a car outside. Respondents distinguish intermittent bursts that interrupt concentration, such as drilling, from steady drones such as traffic hum that establish a persistent background. Building-borne transmission through footsteps, plumbing and vibrating walls is described as particularly difficult to shut out, because the sound arrives through the structure rather than through a window that can be closed.

4.6 Recurrence and the sleep–vigilance pathway

The modal account is not of an extraordinary one-off event but of ordinary domestic life made difficult by repetition. Eighty-five of the 106 neighbour episodes (80.2 percent) are coded as recurring patterns, and 39 of the 43 exposed sessions contain at least one. Night-time exposure appears in 35 sessions. Respondents repeatedly distinguish sounds acceptable during the day from the same sounds at night, when they violate an expectation of rest and are more salient against a quieter background. S28 put the norm plainly: "Night-time means silence. So noises at night seem much louder to me than noises during the day."

Sleep disruption appears in 24 of the 43 neighbour-episode sessions and stress or hypervigilance in 25. These narratives give content to the within-person coefficients in Table 2. Y37 describes recurring parties and fights:

> *"You go to bed early and want to sleep, but they are partying and fighting and do not let you sleep. Suddenly it is three or five in the morning and there it is: noise, noise, noise."* (Y37, translated from Spanish)

The accounts also describe how a short acoustic event becomes a long waking one. Respondents report being jolted awake and then remaining alert, mentally rehearsing the next day's obligations or listening for whether the disturbance will resume. A sequence of such nights carries into the day as fatigue, irritability and reduced patience. Some respondents describe a further stage, in which the body prepares for a disturbance that has not yet occurred, tensing around predictable peak hours even on evenings that stay quiet. The burden then begins before the sound does.

At the same time, the interviews reveal why a binary distinction between persistence and recovery is crude. "Getting used to" noise can denote at least three things: a reduced attentional or physiological response, a practical adaptation, or resignation because the source cannot be changed. AJ12 describes becoming accustomed to traffic and building noise after initially waking easily. S25 reports little personal effect even when a sister could not sleep. AJ7 uses habituation and resignation almost interchangeably. As Section 2.1 noted, subjective adaptation and physiological adaptation are not the same thing, so an account of having got used to a neighbour cannot be read as evidence that the exposure has ceased to be costly.

4.7 Noise as social information

Respondents do not respond to loudness alone. The same sound is tolerated or resented depending on whether it is read as necessary, or as discretionary and imposed. Daytime clatter from work is generally granted legitimacy even when it is jarring; balcony music at 23:00 is not, because its time, place or volume could in principle have been otherwise.

The attribution appears in some cases to be manufactured by the exposure rather than read off it. Episodes triggered by everyday domestic sources—footfall, household appliances, televisions and voices—account for fifty episodes across 31 sessions, of which thirty-six record annoyance or stress and eight record social conflict. Not all of these are innocuous: some voice and audio episodes involve shouting, door-banging or deliberately loud music. But many describe activity nobody could reasonably have avoided—a washing machine at one in the morning, a neighbour upstairs in heels, a radio on a Sunday, friends talking in a flat. Conflict attaches to neighbours rather than to impersonal sources: among the 37 sessions in which both neighbour-linked and traffic-linked episodes are observed, 17 record conflict attached to a neighbour but not to traffic, while one records the reverse (exact McNemar $p = 0.00015$). One respondent, describing an evening with visitors, captures the asymmetry from both sides: they felt they were talking normally, and for the neighbours it was intolerable. Hearing a neighbour through a wall is itself sufficient to generate an attribution of intent, and the attribution can then become a dispute.

Whether a disturbance is bounded also shapes the response. Respondents describe open-ended nights as qualitatively worse than long but predictable ones, because not knowing when a noise will stop sustains monitoring after the sound itself has become tolerable.

The clearest case of meaning rather than acoustics driving the response lies outside the neighbour-episode definition. Respondents consistently read fireworks and gunshots as signals of the arrival or presence of drug trading rather than as distinct acoustic events, and we group them for reporting on that basis while the underlying coded dataset retains them as separate categories. Nine of the 19 episodes in this combined category (47.4 percent) carry fear or insecurity, the largest share among the publication trigger groupings. The same physical sound would carry no such burden where it signified a celebration.

4.8 Control, constraint and the shrinking refuge

What makes a disturbance feel chronic is often the combination of recurrence, anticipation and the absence of an effective remedy. Thirteen of the 43 neighbour-episode sessions contain some form of constraint—fear of confronting the source, a perceived social norm against confrontation, or an institution that fails to provide lasting relief. Nineteen contain some evidence of agency, either direct confrontation or an institution that appears effective. The same respondent may experience both in different circumstances.

S26 describes a repeated sequence in which a complaint or police visit suppresses the disturbance only temporarily: “That stops it, but afterward it happens again, and so on—like a cycle.” AJ18 makes the anticipatory component explicit:

> *“It was upsetting because I already knew it was coming. And we could do nothing [. . .] but it was constant. It is very upsetting to experience it so many days each week.”* (AJ18, translated from Spanish)

Between direct confrontation and institutional recourse lies a wide repertoire of private adjustment. Respondents close windows, rearrange rooms, move a bedroom away from the street, mask sound with a fan or television, and shift their routines around predictable disturbances. These provide partial relief while redistributing the burden: closing windows costs thermal comfort in summer, masking introduces a new constant hum, and moving a child's bedroom crowds another room. Where these adjustments accumulate without resolving the source, respondents describe remaining physically at home while feeling mentally on guard, with fewer spaces and hours that feel fully their own.

Housing structure determines how much of this is possible. Y37 links lack of control to the resources required to improve insulation; S28 describes a subdivided dwelling in which every alarm awakened all occupants; AJ18 describes being effectively obliged to hear footsteps, stairs and music. The housing environment therefore helps determine transmission, while perceived safety, trust and satisfaction may be consequences as well as antecedents. This supports the decision not to use contemporaneous residential satisfaction or a broad subjective context index as ordinary controls in the preferred quantitative specification.

Moving intentions have the same problem of ordering. In the panel, stronger plans to move are closely associated with reported neighbour noise. In the interviews, AJ6 describes the direction of the relationship without ambiguity, clearly show why the desire to move, and especially wanting to move but being unable to, cannot be assumed to be a clean pre-exposure measure. In some cases it is part of the outcome being studied:

> *"It had become a family issue that affected everyone. That was when we began thinking about moving house. [. . .] You start to see that quality of life also begins with having a quieter space."* (AJ6, translated from Spanish)

4.9 Heterogeneous response without fixed "types"

The quantitative interaction analysis found few robust observable moderators of the within-person relationship. The interviews help explain this result: sensitivity is often situational rather than a stable characteristic recoverable from age, sex, education or income. It changes with the time of day, the activity being interrupted, the relationship with the source, and what the sound is taken to mean.

Table 6 sets out the case strata used to examine this variation. The 14-case matrix deliberately includes the modal sleep-and-stress pattern, resilient accounts, low-survey/high-narrative and high-survey/muted-narrative contrasts, constrained and agentic responses, and negative cases at the boundary of the noise construct.

Within-household contrasts are especially revealing. In S25, the interviewed respondent describes calling the police because a sister could not sleep, but then adds:

> *"It was because my sister did not sleep well. [. . .] But it never affected me personally; it never caused me a problem."* (S25, translated from Spanish)

S29 similarly identifies a mother as the more sensitive member of the household, while S23 reports being habituated even though a daughter is readily distracted when studying. A household-level noise report can therefore coexist with very different person-level outcomes. This is a plausible source of attenuation in both the cross-sectional and within-person estimates,

Table 6: Purposefully selected qualitative case patterns and their analytic contribution

Pattern	Illustrative cases	Contribution to interpretation
Chronic adverse exposure with constrained response	Y37, AJ7, S26	Recurring night noise combines broken sleep and stress with poor insulation, resignation, fear and remedies that provide only temporary relief.
Low survey score but strong narrative response	AJ12, S28	Waking, vibration and night-time disturbance are described despite neighbour-nuisance scores below the high threshold, suggesting timing, adaptation or reference-period mismatch.
Household spillover and within-household heterogeneity	AJ6, S25, S29, S23	Children, siblings and parents respond differently to the same dwelling environment; a single household respondent can therefore conceal effects on another member.
Agency and context-dependent tolerance	AJ44, MP20	Direct confrontation can stop a disturbance, while a good relationship with the source can make the same physical sound more tolerable.
Anticipatory or "trapped" response	AJ18	The burden begins before the sound when repeated exposure is expected and the respondent sees no effective action other than repeatedly calling the police.
Survey–narrative construct boundary	Y39, S23	High closed-form nuisance can coexist with muted personal response or with sleep disturbance attributed to gunshots rather than a strict-neighbour episode.
Group-session boundary	AJ45	Co-residents distinguish neutral ambient sound from night-time television and uncontrolled parties, but merged speaker attribution limits person-specific inference.

Notes: The complete case matrix contains 14 purposively selected interview/household sessions and 29 transcript-verified extracts. The cases were selected for maximum analytic variation and negative-case analysis rather than statistical representativeness. AJ45 is a three-person household interview; it remains one analytic cluster.

particularly when the survey respondent is not the household member whose sleep is most affected.

Appraisal also depends on the relationship with the source. MP20 tolerates an adolescent neighbour practising the drums because the families get along well and the activity is interpreted as learning. The same respondent finds birdsong in the morning pleasant and cats on the roof at night frightening, because in the early hours she cannot tell whether what she hears above her is cats or people. Y37 is understanding about the ordinary activities of an elderly neighbour but angry about late parties and fighting from another household. AJ44 reports anger at late music but also describes a successful direct request to turn it off. Much of the relevant heterogeneity is therefore relational and time-varying, and poorly represented by standard socioeconomic characteristics.

4.10 Concordance, discordance and negative cases

The closed-form questionnaire and the narratives correspond on average, but not perfectly. Among the 21 sessions in which the maximum closed-form nuisance score for neighbour voices, music or footfall is at least four on the five-point scale, 71.4 percent contain a qualitative sleep pathway or outcome, compared with 37.5 percent of the 24 lower-scoring sessions (two-sided Fisher exact $p = 0.036$). The corresponding stress rates are 71.4 and 41.7 percent ($p = 0.071$). With only 45 sessions these comparisons are exploratory, but higher reported nuisance is associated with the mechanisms the panel results would lead us to expect.

The exceptions are equally useful. AJ12 and S28 describe substantial waking, vibration or night-time disturbance despite closed-form scores below the high threshold; in both cases adaptation or a previous dwelling complicates temporal alignment. Y39 reports high nuisance

and frequent sleep disruption, but the waking episode is attributed to gunshots or fireworks and is therefore not coded as a neighbour sleep effect. S23 reports high neighbour nuisance while describing personal habituation and little sleep effect, although a daughter has difficulty concentrating. These reveal differences in reference period, source definition, current versus past dwelling, household respondent, and the distinction between noticing a sound and suffering a particular outcome. The negative cases also discipline causal interpretation; circumstances can alter sensitivity. Y37 reports that noise was hardest to bear while studying, and several respondents describe competing demands lowering their tolerance. Outcome-driven perception is therefore possible in individual episodes.

Taken togther, however, the qualitative interviews provide a nuanced and vivid body of support for a noise-moderated interpretation of the quantitative patterns in the survey data. No single reverse-causality story explains the collection of patterns. Respondents describe concrete sources, night-time timing, repeated failed remedies, household members responding differently to the same source, and relief when the source changes or stops. Combined with the quantitative finding that prior outcome levels and changes generally do not predict onset, the narratives make a wholly outcome-generated reporting explanation less plausible.

4.11 Health statements and the boundary of the evidence

Direct health statements are rare. Only two neighbour-episode sessions contain a coded self-described health symptom, and the relevant passages concern next-day fatigue or headache after poor sleep. They are not diagnoses. The interviews therefore support the plausibility of sleep disruption, stress and hypervigilance as mechanisms, but they cannot confirm the exploratory heart-disease result in Table 4. This boundary is useful: the qualitative material strengthens the interpretation of the well-being, sleep and mental-health findings much more directly than it strengthens the diagnosis analysis.

The qualitative strand therefore bears most directly on the well-being, sleep and mental-health findings. It identifies recurrent night-time disturbance, anticipation and low control as plausible mechanisms, while showing why benign appraisal and substantial burden can coexist within the same person or household.

5 Synthesis: what the two evidence bases show together

The two bodies of evidence address different parts of the same question. The longitudinal analysis shows whether well-being and health move with reported neighbour noise within the same person, and whether observed selection or prior deterioration accounts for that movement. The interviews show what changes in reported neighbour noise correspond to in daily life: what enters the home, when and how often it occurs, how residents interpret it, and whether they can bring it to an end. Read together, the evidence identifies neighbour noise as a residential stressor produced by three linked processes: transmission through the dwelling, appraisal of the source and situation, and control over whether the disturbance continues. The value of combining the two studies lies not simply in finding similar patterns in different data. Each supplies information the other cannot. The panel establishes the consistency and temporal structure of the association across a national urban sample; the interviews identify the physical, social and behavioural processes represented by that association.

5.1 A within-person relationship with substantive content

The central longitudinal result is that neighbour noise is associated with changes in outcomes within the same respondent, not only with differences between people. Life satisfaction is lower, and sleep problems and depressive symptoms are higher, in waves when a respondent reports more neighbour noise. These relationships remain remarkably stable as changing income, subjective status, home ownership and residential history are added. They also survive fixed-effects, first-difference, non-linear and alternative-threshold specifications.

The timing evidence strengthens this interpretation. Prior levels and recent changes in life satisfaction, sleep and depressive symptoms do not generally predict the onset of high reported neighbour noise, and next-wave noise does not predict current outcomes once current noise is taken into account. Residential circumstances matter for who begins reporting noise, but a simple account in which worsening well-being itself generates a new noise report does not fit the overall temporal pattern.

The interviews show what this within-person movement can represent. Accounts are anchored in identifiable sources and changes in residential conditions: a neighbouring household begins holding late parties, a new source becomes audible through a shared structure, a recurring pattern intensifies, or a source moves away. Respondents describe deterioration while the disturbance continues and relief when direct action succeeds or the source stops. Reported neighbour noise therefore carries substantive information about changes in the home environment and the burden those changes impose.

The larger between-person associations are consistent with the same account operating over longer periods. Some residents describe repeated night-time disturbance, highly permeable dwellings and remedies that work only briefly. Their accounts make cumulative and persistent burden plausible, while the within-person estimates provide the paper's clearest evidence that the relationship is not simply a fixed difference between people who live in different places.

5.2 How sound becomes a residential stressor

Sleep provides the most direct bridge between the two studies. Sleep problems rise consistently with reported neighbour noise in the panel. In the interviews, residents describe delayed sleep, repeated waking, listening for the next sound and carrying fatigue and irritability into the following day. The problem is seldom a single exceptional event. It is the repetition of ordinary and extraordinary sounds across nights that turns an acoustic interruption into an enduring residential burden.

The interviews also show why the burden extends beyond the period in which a sound is audible. When a disturbance recurs at a recognisable hour, residents begin to anticipate it. They monitor whether a party is starting, whether a vehicle will return or whether a temporarily quiet neighbour will resume. The body and attention prepare before the sound arrives, and alertness can continue after it ends. This anticipatory process gives concrete meaning to the relationship between reported neighbour noise and depressive symptoms: the stressor occupies time, attention and emotional space beyond the acoustic event itself.

Appraisal shapes this process without making it arbitrary. Residents respond to timing, recurrence, predictability, perceived discretion and their relationship with the source. Daytime activity understood as necessary can be tolerated even when conspicuous, while an avoidable disturbance late at night is experienced as an intrusion. A friendly relationship can make a

neighbour's activity easier to accommodate; repeated disregard can make the same kind of sound increasingly difficult to bear. The fact that respondents sometimes describe both benign and adverse reactions is therefore not an inconsistency; it shows that neighbour noise is experienced within a social and temporal setting.

Control determines whether this stressor is brief or sustained. Successful direct requests and effective institutional action can restore quiet. Fear of confrontation, weak enforcement and remedies that last only a few days produce a recurring cycle of disturbance, complaint and relapse. Private adaptations—closing windows, rearranging rooms, masking sound or changing routines—can reduce immediate exposure but also reduce thermal comfort, space, autonomy or rest. As control narrows, the home becomes a less reliable refuge even during periods of temporary quiet.

Taken together, these processes explain why the three principal outcomes move in the same adverse direction. Sleep records the immediate interruption; depressive symptoms capture part of the accumulating emotional and cognitive burden; and life satisfaction reflects the wider loss of comfort, autonomy and security within the home.

5.3 Residential context as both origin and consequence

The combined evidence also clarifies the role of residential context. Neighbour noise is not randomly distributed across dwellings or neighbourhoods. Prior residential conditions and other neighbourhood problems predict the onset of high reported noise, while the interviews identify building permeability, proximity, dwelling subdivision and changes in nearby occupancy as conditions that determine what residents hear. Residential context therefore helps generate exposure.

At the same time, neighbour noise reshapes the context in which it occurs. In the panel, reported noise moves closely with lower home satisfaction, lower perceived safety, lower neighbour trust and stronger intentions to move. The interviews show the direction of this process in individual cases: repeated disturbance damages relations with neighbours, makes residents less secure about confronting the source and leads households to discuss leaving a dwelling they had previously expected to remain in. Residential context is therefore both antecedent and outcome. Contemporaneous trust, satisfaction and moving intentions often form part of the burden itself.

This point is especially important for mobility. Moving may end a particular exposure, but the decision to consider leaving is already evidence that the home is no longer functioning as expected. Relocation also imposes financial, social and practical costs, and is not equally available to all households. Wanting to move but being unable to do so consequently captures more than residential preference: it can mark a loss of control over both the source and the possibility of escape.

5.4 Heterogeneity without fixed noise-sensitive types

The two bodies of evidence also converge on the nature of heterogeneity. The quantitative panel analysis provides little evidence that responses divide into stable groups that can be identified by age, sex, education, income or conventional indicators of socioeconomic position. The qualitative interviews show why: reactions vary with the time of day, the activity being interrupted, the relationship with the source, the meaning attached to the sound and the availability of an effective response.

The distributional implications likewise extend beyond who encounters a noisy neighbour. Conventional socioeconomic characteristics differ little across persistent-low, changing and persistent-high noise histories. Inequality may instead arise in what happens after noise begins: the acoustic protection provided by the dwelling, the safety and social cost of confronting the source, the responsiveness of landlords or public institutions, and the resources required to insulate or move. The capacity to secure relief may therefore constitute a central dimension of the burden.

5.5 Integrated interpretation

The two studies support a coherent account of neighbour noise. Sound crosses a residential boundary; timing, recurrence and social meaning determine how intrusive it becomes; and control determines whether the problem is resolved or becomes an anticipated feature of domestic life. Across the panel, this process is associated with deterioration in well-being, sleep and depressive symptoms within the same people. Across the interviews, it appears in recurring night-time disturbance, waking, vigilance, strained relations, costly adaptation and relief when effective action becomes possible.

Residential selection is part of this account, but it does not explain the whole pattern. Nor are reactions reducible to a fixed disposition to be sensitive to sound. The evidence instead points to an interaction among buildings, situations, social relationships and the capacity to act. This is the paper's central mixed-method conclusion: neighbour noise is an ordinary residential exposure that can erode sleep, well-being and the security of the home when transmission is high and effective relief is unavailable.

The exploratory diagnosis results raise a separate longer-run question for future research. The integrated evidence developed here bears most directly on life satisfaction, sleep and depressive symptoms, for which the longitudinal and qualitative findings speak to the same underlying process.

6 Concluding discussion

Set against the existing literature, the quantitative contribution here is confirmatory. Earlier analysis of Dutch panel data established the within-person relationship and addressed reverse causation by restricting the sample to initially healthy respondents (Weinhold 2015); subsequent work established temporal ordering over longer intervals and in both directions (Beutel et al. 2020, Cassens et al. 2025). We find the same patterns in a different country from a different survey, and subject them to a wider range of selection and sensitivity checks. A distinctive contribution is our qualitative fieldwork, its structured episode coding and analysis, and the synthesis of quantitative and qualitative results into a coherent narrative.

6.1 Housing and policy implications

The policy implications follow from the same three processes that organise the evidence: transmission, appraisal and control. Neighbour noise cannot be addressed only as individual incivility, because the dwelling determines how much ordinary domestic activity crosses into another home. Nor can it be addressed only through construction, because timing, recurrence, social relations

and the availability of a remedy determine whether audible sound becomes a sustained residential stressor. Effective policy must act on all three margins.

At the transmission stage, stronger acoustic standards can reduce the number of situations in which ordinary domestic life becomes involuntarily shared across walls, floors and ceilings. Recent evidence that complaints fall sharply under stricter floor-impact-sound regimes suggests that building regulation can change lived experience as well as measured acoustic performance (Shin et al. 2026). Acoustic labelling and disclosure could also make an otherwise hidden dwelling characteristic more visible before a purchase or rental (Şentop Dümen & Rasmussen 2025*b*). For the existing housing stock, where rebuilding is impossible, targeted retrofit, clearer assignment of responsibility between owners and housing providers, and attention to subdivision and building services may offer more realistic routes to reducing transmission.

The social character of neighbour noise requires a complementary institutional response. Dense residential life necessarily involves audible activity, and the interviews contain many examples of sounds regarded as reasonable, necessary or harmless. The relevant distinction is not between perfect silence and any audible neighbour, but between ordinary domestic sound and disturbance that becomes recurrent, avoidable, threatening or concentrated at times when residents reasonably expect to sleep. Clear quiet hours, building or tenancy rules and proportionate standards can help define that boundary without requiring every affected resident to negotiate it from scratch.

Control is especially important because direct bargaining is often costly. Neighbours may disagree about what occurred, whether it was avoidable and whose use of the home should prevail. A complaint can also damage the relationship on which an informal solution depends (Coase 1960, Stokoe & Hepburn 2005). A graduated system of response is therefore preferable to either non-intervention or immediate punitive enforcement: accessible advice, documented notification, mediation where appropriate, involvement of landlords or building managers, and enforceable remedies for repeated breaches. Such a system should allow residents to seek relief without forcing them into repeated personal confrontation or requiring them to recast every problem as interpersonal hostility.

A remedy framework built on identifying a wrongdoer also misdescribes much of the problem. A substantial share of what residents report arises from activity nobody could reasonably have avoided: a washing machine at night, footfall overhead, a conversation in an adjacent flat. Where a building transmits ordinary domestic life, both households bear a cost. One hears what it would rather not; the other must curtail ordinary use of its home or be cast as inconsiderate. Enforcement aimed at the producer sits awkwardly with this, because it assigns fault where the binding constraint is the structure, and it helps explain why residents so often describe complaining as socially expensive: to complain is to accuse. Remedies that act on transmission, or that provide a route to relief which does not require one neighbour to name another, avoid manufacturing a dispute out of a construction defect.

This institutional gap is clearly visible in the Chilean interviews, but is hardly unique—internationally, a weak regulatory environment is the norm (Hammer et al. 2018). Ordinary domestic activity falls largely outside the national environmental-noise standard and is instead dispersed across municipal rules, co-ownership arrangements and police response (Ministerio del Medio Ambiente, Gobierno de Chile 2011, Congreso Nacional de Chile 2022). The result is that the availability of relief can depend heavily on dwelling type, tenure and local institutional capacity. A clearer and more consistent route for handling recurring residential disturbance would complement rather than replace building standards.

Mobility should not be treated as the default remedy. Moving may end a particular exposure, but considering exit is itself evidence that the dwelling is no longer providing the expected security and refuge of a home. Relocation also carries financial, social and practical costs and is least available to households with the fewest alternatives. The strong relationships with plans to move and wanting to move while unable therefore capture an important dimension of the burden rather than an easy route by which the problem solves itself.

6.2 Scope and research agenda

The strongest conclusions concern life satisfaction, sleep and depressive symptoms. These outcomes move consistently with reported neighbour noise within respondents, and the interviews describe the same sleep, vigilance and residential-control processes in daily life. The diagnosis analysis asks whether these pathways may also have consequences over longer horizons. The concentration of incident heart disease among older respondents with high early exposure is notable and resembles the earlier Netherlands result (Weinhold 2015), but the number of diagnoses is small and the hypertension results do not show the same pattern. The physical-health finding is therefore interpreted as an addition to the observational hypothesis-generating evidence base that motivates replication with a larger, purpose-built research design.

The present design estimates the average within-person relationship across annual survey waves. It cannot reconstruct the precise onset, duration and recovery path of individual disturbances. The panel also lacks acoustic measurements, detailed building characteristics and information about the specific source or neighbour associated with each report. Respondent fixed effects remove stable differences among people and their long-run environments, while survey-wave effects remove common temporal shocks; local changes that simultaneously affect neighbour noise and well-being can still matter. Because reported neighbour noise can refer to another household member while the outcomes are respondent-specific, genuine household burden may also be imperfectly aligned with the individual outcome recorded in a given wave.

The interviews address many of these interpretive gaps but have a different scope. They were purposively recruited in Temuco, are not linked to the panel respondents and are designed to identify processes and contrasts rather than population prevalence. Their contribution is to show how neighbour sound enters domestic life, why it produces different reactions across settings and household members, and how residents experience agency or constraint. The fact that the national panel and the Temuco interviews are independent is a strength for comparison: the recurring patterns of sleep disruption, anticipation and limited control were not produced by interviewing a selected subset of the survey respondents.

The next generation of research should combine the strengths of both approaches. Repeated source-specific reports collected at shorter intervals could establish when a disturbance begins, whether it persists and how quickly outcomes recover when it stops. Acoustic measurements and dwelling characteristics could identify the contribution of physical transmission, while parallel reports from multiple household members could reveal how a shared environment produces unequal individual consequences. Information on neighbour turnover, complaints, building remediation and institutional response would allow researchers to distinguish changes in sound from changes in control. Designs exploiting externally timed occupancy changes, retrofitting programmes or enforcement reforms would be especially valuable for estimating causal effects and comparing the effectiveness of alternative remedies.

Instruments should also record both sides of the exposure. Existing surveys, including the one used here, ask only whether a respondent is disturbed. None asks whether a respondent has

been the subject of a complaint, or has curtailed ordinary activity in order to avoid one. That omission makes the burden appear one-sided when it plausibly falls on both parties, and it leaves the producing household's experience entirely unmeasured. Measuring both roles would also allow a direct test of whether transmission, rather than inconsiderate behaviour, is the binding constraint.

Neighbour noise is an ordinary exposure with unusually intimate consequences. It enters through the fabric of the dwelling, acquires meaning through relations between neighbours, and becomes persistent when residents lack a safe and effective route to relief. Across seven waves of the Chilean panel, life satisfaction, sleep and depressive symptoms move adversely with the same respondent's reported neighbour noise. The Temuco interviews show how recurring nights, anticipation, social meaning and unsuccessful attempts to restore quiet turn that exposure into a broader loss of residential well-being.

Together, the findings place neighbour noise at the intersection of housing quality, public health and local governance. A quiet home is not simply a matter of individual tolerance or neighbourly goodwill. It is a basic residential amenity jointly produced by buildings, shared norms and institutions capable of providing effective relief.

References

Ashwin, J., Chhabra, A. & Rao, V. (2026), 'Using large language models for qualitative analysis can introduce serious bias', *Sociological Methods & Research* **55**(3), 795–839. First published online in 2025.

Babisch, W. (2014), 'Updated exposure-response relationship between road traffic noise and coronary heart diseases : A meta-analysis', *Noise & Health* **16**(68), 1–9.

Bayo Jimenez, M. T., Gericke, A., Frenis, K., Rajlic, S., Kvandova, M., Kröller-Schön, S., Oelze, M., Kuntic, M., Kuntic, I., Daiber, A. & Münzel, T. (2023), 'Effects of aircraft noise cessation on blood pressure, cardio- and cerebrovascular endothelial function, oxidative stress, and inflammation in an experimental animal model', *Science of the Total Environment* **903**, 166106.

Beghelli, S., De Coulon, A. & O'Mahony, M. (2023), 'Health benefits of reducing aircraft pollution: Evidence from changes in flight paths', *Journal of Population Economics* **36**, 2691–2712.

Beutel, M. E., Brähler, E., Ernst, M., Klein, E., Reiner, I., Wiltink, J., Michal, M., Wild, P. S., Schulz, A., Münzel, T., Hahad, O., König, J., Lackner, K. J., Pfeiffer, N. & Tibubos, A. N. (2020), 'Noise annoyance predicts symptoms of depression, anxiety and sleep disturbance 5 years later: Findings from the gutenberg health study', *European Journal of Public Health* **30**(3), 516–521.

Boes, S., Nüesch, S. & Stillman, S. (2013), 'Aircraft noise, health, and residential sorting: Evidence from two quasi-experiments', *Health Economics* **22**(9), 1037–1051.

Braun, V. & Clarke, V. (2006), 'Using thematic analysis in psychology', *Qualitative Research in Psychology* **3**(2), 77–101.

Braun, V. & Clarke, V. (2019), 'Reflecting on reflexive thematic analysis', *Qualitative Research in Sport, Exercise and Health* **11**(4), 589–597.

Cassens, L., van Cranenburgh, S., Calvert, S. & Kroesen, M. (2025), 'Longitudinal analysis of the influence of personality on noise annoyance and bidirectional effects between noise and health outcomes', *Noise & Health* **27**(124), 80–88.

Centro de Estudios de Conflicto y Cohesión Social (COES) (2025), 'Estudio longitudinal social de chile (ELSOC): Ondas 1–7, 2016–2023', Panel data and methodological documentation. Accessed 15 August 2026.
URL: *https://conferencias.coes.cl/encuesta-panel/*

Chen, Y., Hansell, A. L., Clark, S. N. & Cai, Y. S. (2023), 'Environmental noise and health in low-middle-income-countries: A systematic review of epidemiological evidence', *Environmental Pollution* **316**, 120605.

Chernysheva, L. (2024), 'The vagaries of sonic neighborly life: Privacy at home and sound politics in post-socialist mass housing', *Space and Culture* **27**(2), 288–301.

Coase, R. H. (1960), 'The Problem of Social Cost', *The Journal of Law and Economics* **3**(1), 1.

Congreso Nacional de Chile (2022), 'Ley 21.442: Nueva ley de copropiedad inmobiliaria', Diario Oficial de la República de Chile.

Covassin, N. & Singh, P. (2016), 'Sleep duration and cardiovascular disease risk: Epidemiologic and experimental evidence', *Sleep Medicine Clinics* .

Dai, S.-C., Xiong, A. & Ku, L.-W. (2023), LLM-in-the-loop: Leveraging large language model for thematic analysis, *in* 'Findings of the Association for Computational Linguistics: EMNLP 2023', Association for Computational Linguistics, Singapore, pp. 9993–10001.

Diegert, P., Masten, M. A. & Poirier, A. (2026), 'Assessing omitted variable bias when the controls are endogenous', *arXiv preprint arXiv:2206.02303* . Latest revision; first circulated in 2022.

Fan, Y. & Weinhold, D. (2022), 'Urban noise, sleep disruption and health', *Applied Economics* **54**(50), 5782–5799.

Frenis, K., Kalinovic, S., Ernst, B. P., Kvandova, M., Al Zuabi, A., Kuntic, M., Oelze, M., Düerr, G. D., Bayo Jimenez, M. T., Steven, S., Daiber, A. & Münzel, T. (2021), 'Long-term effects of aircraft noise exposure on vascular oxidative stress, endothelial function and blood pressure: No evidence for adaptation or tolerance development', *Frontiers in Molecular Biosciences* **8**, 814921.

Gangwisch, J. E. (2014), 'A review of evidence for the link between sleep duration and hypertension', *American journal of hypertension* **27**(10), 1235–1242.

Hahad, O., Kuntic, M., Al-Kindi, S., Kuntic, I., Gilan, D., Petrowski, K., Daiber, A. & Münzel, T. (2025), 'Noise and mental health: Evidence, mechanisms, and consequences', *Journal of Exposure Science & Environmental Epidemiology* **35**, 16–23.

Hamilton, L., Elliott, D., Quick, A., Smith, S. & Choplin, V. (2023), 'Exploring the use of AI in qualitative analysis: A comparative study of guaranteed income data', *International Journal of Qualitative Methods* **22**, 16094069231201504.

Hammer, M. S., Fan, Y., Hammer, S. S., Swinburn, T. K., Weber, M., Weinhold, D. & Neitzel, R. L. (2018), 'Applying a novel environmental health framework theory (I-ACT) to noise pollution policies in the united states, united kingdom, and the netherlands', *Journal of Environmental Planning and Management* **61**(12), 2111–2132.

Hennink, M. M., Kaiser, B. N. & Marconi, V. C. (2017), 'Code saturation versus meaning saturation: How many interviews are enough?', *Qualitative Health Research* **27**(4), 591–608.

Hill, C., Dahil, A., Simpson, G., Hardisty, D., Keast, J., Pinn, C. K. & Dambha-Miller, H. (2026), 'Large language models for thematic analysis in healthcare research: A blinded mixed-methods comparison with human analysts', *PLOS Digital Health* **5**(4), e0001189.

Jensen, H. A. R., Rasmussen, B. & Ekholm, O. (2019), 'Neighbour noise annoyance is associated with various mental and physical health symptoms: Results from a nationwide study among individuals living in multi-storey housing', *BMC Public Health* **19**, 1508.

Jones, K. M. L. (2025), 'Generative AI in qualitative research and related transparency problems: A novel heuristic for disclosing uses of AI', *International Journal of Qualitative Methods* **24**, 16094069251404329.

Karjus, A. (2025), 'Machine-assisted quantitizing designs: Augmenting humanities and social sciences with artificial intelligence', *Humanities and Social Sciences Communications* **12**, 277.

Knutson, K. L., Ryden, A. M., Mander, B. A. & Van Cauter, E. (2006), 'Role of sleep duration and quality in the risk and severity of type 2 diabetes mellitus', *Archives of internal medicine* **166**(16), 1768–1774.

Kröller-Schön, S., Daiber, A., Steven, S., Oelze, M., Frenis, K., Kalinovic, S., Heimann, A., Schmidt, F. P., Pinto, A., Kvandova, M., Vujačić-Mirski, K., Filippou, K., Dudek, M., Bosmann, M., Klein, M., Bopp, T., Hahad, O., Wild, P. S., Frauenknecht, K., Methner, A., Schmidt, E. R., Rapp, S., Mollnau, H. & Münzel, T. (2018), 'Crucial role for Nox2 and sleep deprivation in aircraft noise-induced vascular and cerebral oxidative stress, inflammation, and gene regulation', *European Heart Journal* **39**(38), 3528–3539.

Lazarus, R. S. & Folkman, S. (1984), *Stress, Appraisal, and Coping*, Springer, New York, NY.

Leventhall, H. G. (2004), 'Low frequency noise and annoyance.', *Noise & health* **6**(23), 59–72.

Méndez, M. L. & Otero, G. (2018), 'Neighbourhood conflicts, socio-spatial inequalities, and residential stigmatisation in santiago, chile', *Cities* **74**, 75–82.

Miles, M. B., Huberman, A. M. & Saldaña, J. (2014), *Qualitative Data Analysis: A Methods Sourcebook*, 3 edn, SAGE, Thousand Oaks, CA.

Ministerio del Medio Ambiente, Gobierno de Chile (2011), 'Decreto supremo n° 38/2011: Norma de emisión de ruidos generados por fuentes que indica', Diario Oficial de la República de Chile.

Morgan, D. L. (2023), 'Exploring the use of artificial intelligence for qualitative data analysis: The case of ChatGPT', *International Journal of Qualitative Methods* **22**, 16094069231211248.

Morgan, D. & Tsai, S. C. (2015), 'Sleep and the endocrine system', *Critical care clinics* **31**(3), 403–418.

Münzel, T., Molitor, M., Kuntic, M., Hahad, O., Röösli, M., Engelmann, N., Basner, M., Daiber, A. & Sørensen, M. (2024), 'Transportation noise pollution and cardiovascular health', *Circulation Research* **134**(9), 1113–1135.

Münzel, T., Sørensen, M. & Daiber, A. (2021), 'Transportation noise pollution and cardiovascular disease', *Nature Reviews Cardiology* **18**(9), 619–636.

Niemann, H., Bonnefoy, X., Braubach, M., Hecht, K., Maschke, C., Rodrigues, C. & Röbbel, N. (2006), 'Noise-induced annoyance and morbidity results from the pan-European LARES study.', *Noise & health* **8**(31), 63–79.
URL: *http://www.ncbi.nlm.nih.gov/pubmed/17687182*

Niemann, H. & Maschke, C. (2004), Noise effects and morbidity, Technical report, World Health Organization.

Nissenbaum, M. A., Aramini, J. J. & Hanning, C. D. (2012), 'Effects of industrial wind turbine noise on sleep and health', *Noise & Health* **14**(60), 237–243.

Office of the New York State Comptroller (2018), Noise in new york city neighborhoods, Technical report, Office of the State Deputy Comptroller for the City of New York. Accessed 18 August 2026.
URL: *https://www.osc.ny.gov/press/releases/2018/01/dinapoli-silencing-excessive-nyc-noise-major-challenge*

Office of the New York State Comptroller (2025), NYC311 monitoring tool, Report 3-2026, Office of the State Deputy Comptroller for the City of New York. Report number follows the Comptroller's fiscal-year series. Accessed 18 August 2026.
URL: *https://www.osc.ny.gov/files/reports/pdf/report-3-2026.pdf*

Osborne, M. T., Radfar, A., Hassan, M. Z. O., Abohashem, S., Oberfeld, B., Patrich, T., Tung, B., Wang, Y., Ishai, A., Scott, J. A., Shin, L. M., Fayad, Z. A., Koenen, K. C., Rajagopalan, S., Pitman, R. K. & Tawakol, A. (2020), 'A neurobiological mechanism linking transportation noise to cardiovascular disease in humans', *European Heart Journal* **41**(6), 772–782.

Palmer, C. A. & Alfano, C. A. (2016), 'Sleep and emotion regulation: An organizing, integrative review', *Sleep Medicine Reviews* .

Praag, B. V. & Baarsma, B. (2005), 'Using Happiness Surveys to Value Intangibles: The Case of Airport Noise*', *The Economic Journal* **115**(2000), 224–246.
URL: *http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0297.2004.00967.x/full*

Prescott, M. R., Yeager, S., Ham, L., Rivera Saldana, C. D., Serrano, V., Narez, J., Paltin, D., Delgado, J., Moore, D. J. & Montoya, J. (2024), 'Comparing the efficacy and efficiency of human and generative AI: Qualitative thematic analyses', *JMIR AI* **3**, e54482.

Sandelowski, M., Voils, C. I. & Knafl, G. (2009), 'On quantitizing', *Journal of Mixed Methods Research* **3**(3), 208–222.

Schmidt, F. P., Basner, M., Kröger, G., Weck, S., Schnorbus, B., Muttray, A., Sariyar, M., Binder, H., Gori, T., Warnholtz, A. & Münzel, T. (2013), 'Effect of nighttime aircraft noise exposure on endothelial function and stress hormone release in healthy adults', *European Heart Journal* **34**(45), 3508–3514a.

Segerstrom, S. C. & Miller, G. E. (2004), 'Psychological stress and the human immune system: A meta-analytic study of 30 years of inquiry', *Psychological Bulletin* **130**(4), 601–630.

Şentop Dümen, A. & Rasmussen, B. (2025*a*), 'Emotions elicited by neighbour noise and coping process: A social media research based on two newspaper articles', *Building and Environment* **269**, 112432.

Şentop Dümen, A. & Rasmussen, B. (2025*b*), 'Neighbour noise in multi-storey housing: Public experiences, attitudes and the need for acoustic labelling to ensure consumer protection', *Building Acoustics* **32**(1), 25–49.

Shin, H.-k., Park, S. H. & Kim, S. H. (2026), 'The impact of building regulations on neighbor noise complaints in a longitudinal study', *Journal of Building Engineering* **121**, 115575.

Shukla, C. & Basheer, R. (2016), 'Metabolic signals in sleep regulation: recent insights', *Nature and Science of Sleep* **8**, 9.

Smith, M. G., Cordoza, M. & Basner, M. (2022), 'Environmental noise and effects on sleep: An update to the WHO systematic review and meta-analysis', *Environmental Health Perspectives* **130**(7), 076001.

Sørensen, M., Andersen, Z. J., Nordsborg, R. B., Becker, T., Tjønneland, A., Overvad, K. & Raaschou-Nielsen, O. (2013), 'Long-Term Exposure to Road Traffic Noise and Incident Diabetes: A Cohort Study', *Environmental Health Perspectives* **121**(2).

Stallen, P. J. M. (1999), 'A theoretical framework for environmental noise annoyance', *Noise & Health* **1**(3), 69–80.

Stokoe, E. & Hepburn, A. (2005), "you can hear a lot through the walls': Noise formulations in neighbour complaints', *Discourse & Society* **16**(5), 647–673.

Trudeau, C., King, N. & Guastavino, C. (2023), 'Investigating sonic injustice: A review of published research', *Social Science & Medicine* **326**, 115919.

Ureta, S. (2007), 'Noise and the battles for space: Mediated noise and everyday life in a social housing estate in santiago, chile', *Journal of Urban Technology* **14**(3), 103–130.

Wachinger, J., Bärnighausen, K., Schäfer, L. N., Scott, K. & McMahon, S. A. (2025), 'Prompts, pearls, imperfections: Comparing ChatGPT and a human researcher in qualitative data analysis', *Qualitative Health Research* **35**(9), 951–966.

Walsh, M., Arthurson, K. & Levin, I. (2021), 'Obtrusive intimacy in multi-owned housing: Exploring the impacts on residents' health and wellbeing', *Housing, Theory and Society* **38**(5), 614–630.

Wang, S. S., Glied, S., Williams, S., Will, B. & Muennig, P. A. (2022), 'Impact of aeroplane noise on mental and physical health: A quasi-experimental analysis', *BMJ Open* **12**(5), e057209.

Weinhold, D. (2013), 'The happiness-reducing costs of noise pollution', *Journal of Regional Science* **53**(2), 292–303.

Weinhold, D. (2014), Noise, *in* T. C. Haab & J. C. Whitehead, eds, 'Environmental and Natural Resource Economics: An Encyclopedia', Greenwood, Santa Barbara, CA, pp. 228–231.

Weinhold, D. (2015), Sick of noise: The health effects of loud neighbours and urban din, GRI Working Paper 213, Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science. RePEc handle lsg:lsgwps:wp213. Revised version 2018.
URL: *https://www.lse.ac.uk/granthaminstitute/wp-content/uploads/2015/10/Working-Paper-213-Weinhold.pdf*

World Health Organization Regional Office for Europe (2018), *Environmental Noise Guidelines for the European Region*, World Health Organization Regional Office for Europe, Copenhagen.
URL: *https://www.who.int/publications/i/item/9789289053563*

Zee, P. C. & Turek, F. W. (2006), 'Sleep and health: everywhere and in both directions', *Archives of internal medicine* **166**(16), 1686–1688.

Appendix

Appendix A Quantitative robustness checks

Table 7 collects the principal checks on the average within-person relationships. Panel A retains the five-category exposure and varies the estimator or analytic sample. Panel B changes the exposure parameterisation. Its larger coefficients reflect comparisons over wider differences in reported noise rather than a different substantive conclusion. Panel C reports the Diegert–Masten–Poirier sign-breakdown frontier after respondent and wave residualisation. It calibrates a single omitted time-varying factor against the full observed control block while allowing that factor to correlate with the observed controls. Values above one imply that the omitted factor must be more influential than the benchmark controls in both the exposure and outcome equations before a sign reversal enters the identified set. The benchmark row permits an absolute correlation of up to 0.5; the surrounding rows make the sensitivity to that assumption explicit.

Full replication materials are available at https://github.com/dmweinhold/Neighbour_Noise_Chile_Replication.

Table 7: Robustness of the primary within-person associations

Specification	Life satisfaction	Sleep problems	Depressive symptoms	Person-waves
Panel A. Ordinal noise measure: effect of a one-category increase				
Person and wave fixed effects	−0.0286 ($p < 0.001$)	0.0470 ($p < 0.001$)	0.0610 ($p < 0.001$)	19,536–19,545
+ Survey interview mode	−0.0287 ($p < 0.001$)	0.0467 ($p < 0.001$)	0.0606 ($p < 0.001$)	19,536–19,545
Exclude the 2021 telephone wave	−0.0312 ($p < 0.001$)	0.0500 ($p < 0.001$)	0.0639 ($p < 0.001$)	17,809–17,821
Stable observed commune	−0.0272 ($p < 0.001$)	0.0469 ($p < 0.001$)	0.0650 ($p < 0.001$)	16,970–16,981
At least five outcome–noise pairs	−0.0283 ($p = 0.001$)	0.0443 ($p < 0.001$)	0.0570 ($p < 0.001$)	12,492–12,521
Exclude top 1% of exposure leverage	−0.0287 ($p < 0.001$)	0.0444 ($p < 0.001$)	0.0610 ($p < 0.001$)	18,846–18,855
First differences, excluding 2019–2021	−0.0324 ($p = 0.003$)	0.0465 ($p < 0.001$)	0.0626 ($p < 0.001$)	11,422–11,442
Panel B. Alternative exposure parameterisations				
High noise (“sometimes” or more) versus low	−0.0602 ($p = 0.001$)	0.0769 ($p < 0.001$)	0.1122 ($p < 0.001$)	19,536–19,545
Very high noise (“many times” or “always”) versus lower	−0.0507 ($p = 0.031$)	0.1111 ($p < 0.001$)	0.1350 ($p < 0.001$)	19,536–19,545
Highest category versus “never”	−0.1214 ($p = 0.002$)	0.2579 ($p < 0.001$)	0.2853 ($p < 0.001$)	19,536–19,545
Panel C. DMP sign-breakdown ratio, full observed control block				
Maximum control–confounder correlation $\bar{c} = 0$	3.03	7.15	5.93	16,874–16,882
$\bar{c} = 0.25$	1.84	2.60	2.45	16,874–16,882
$\bar{c} = 0.50$ (benchmark)	1.33	1.62	1.57	16,874–16,882
$\bar{c} = 0.75$	1.07	1.19	1.17	16,874–16,882
$\bar{c} = 1.00$	0.96	0.99	0.99	16,874–16,882

Notes: Outcomes are standardised but retain their natural direction: lower life satisfaction and higher sleep or depressive symptoms are adverse. Depressive symptoms exclude the sleep item. Panel A reports coefficients per one-category increase in the five-category noise measure. Except for first differences, specifications include respondent and survey-wave fixed effects. “Stable observed commune” retains respondents whose observed commune does not change. The leverage restriction removes respondents above the 99th percentile of the within-person exposure sum of squares. Panel B changes the exposure parameterisation; its binary and category-indicator coefficients therefore should not be compared numerically with Panel A. Standard errors are clustered by respondent. Panel C applies the Diegert–Masten–Poirier sensitivity analysis to standardised outcomes, binary high noise and the full control block after residualising each with respect to respondent and survey-wave effects. The benchmark controls are log household income, subjective status, home ownership and years in the neighbourhood. Entries are the smallest common relative-strength value, $\bar{r}_X = \bar{r}_Y$, at which the identified set permits a reversal of the estimated sign; $\bar{c}$ is the assumed maximum absolute correlation between the omitted factor and observed controls. A value of 1 therefore means an omitted factor as influential as the complete observed block in both the noise and outcome equations. Panel C is a calibration under the maintained linear model, not a confidence interval or a test of exogeneity (Diegert et al. 2026).

Appendix A.1 Exploratory socioeconomic heterogeneity

Table 8 reports the complete specification family used to examine whether baseline socioeconomic position moderates the within-person association between neighbour noise and the primary outcomes.

Table 8: Baseline socioeconomic position and within-person noise associations: specification family

Specification	Life satisfaction	Sleep problems	Depressive symptoms	Person-waves	Holm-adjusted p (LS / sleep / depression)
Panel A. Baseline log household income interaction, by specification					
Parsimonious respondent FE, no controls	0.0117 ($p = 0.058$)	−0.0191 ($p = 0.020$)	−0.0094 ($p = 0.045$)	19,345–19,354	—
Full CRE, Table 2 sample and controls (headline)	0.0115 ($p = 0.091$)	−0.0189 ($p = 0.038$)	−0.0097 ($p = 0.062$)	16,874–16,882	—
Outcomes at or after first income observation	0.0121 ($p = 0.060$)	−0.0171 ($p = 0.044$)	−0.0100 ($p = 0.038$)	18,431–18,443	—
Income observed in 2016; outcomes 2017 onward	0.0128 ($p = 0.155$)	−0.0377 ($p = 0.003$)	−0.0150 ($p = 0.046$)	8,704	—
Baseline income standardised within entry wave	0.0117 ($p = 0.085$)	−0.0186 ($p = 0.043$)	−0.0098 ($p = 0.060$)	16,874–16,882	—
Exposure as high-versus-low binary	0.0348 ($p = 0.031$)	−0.0287 ($p = 0.182$)	−0.0171 ($p = 0.143$)	16,874–16,882	—
Exposure fully categorical, joint test	Joint $p = 0.071$	Joint $p = 0.131$	Joint $p = 0.286$	16,874–16,882	—
Panel B. Alternative measures of baseline socioeconomic position, headline specification					
Log household income	0.0115 ($p = 0.091$)	−0.0189 ($p = 0.038$)	−0.0097 ($p = 0.062$)	16,874–16,882	0.546 / 0.385 / 0.495
Education	0.0096 ($p = 0.141$)	−0.0178 ($p = 0.048$)	−0.0091 ($p = 0.070$)	16,874–16,882	0.703 / 0.428 / 0.495
Subjective class	0.0208 ($p = 0.001$)	−0.0061 ($p = 0.494$)	−0.0072 ($p = 0.147$)	16,855–16,863	0.016 / 1.000 / 0.703
Subjective status	0.0177 ($p = 0.011$)	0.0017 ($p = 0.856$)	−0.0046 ($p = 0.366$)	16,874–16,882	0.122 / 1.000 / 1.000
Panel C. Implied within-person noise slope at selected baseline-income values, headline specification					
One standard deviation below mean baseline income	−0.0336 ($p < 0.001$)	0.0673 ($p < 0.001$)	0.0453 ($p < 0.001$)	16,874–16,882	—
One standard deviation above mean baseline income	−0.0106 ($p = 0.261$)	0.0295 ($p = 0.020$)	0.0259 ($p < 0.001$)	16,874–16,882	—

Notes: Entries in Panels A and B are interactions between within-person neighbour-noise exposure and the indicated time-invariant baseline characteristic; two-sided clustered p-values appear beneath the coefficients. Outcomes remain in their original units and natural directions, so positive interactions for life satisfaction and negative interactions for sleep or depressive symptoms indicate attenuation at higher socioeconomic position. Baseline measures are respondents' first observed values and are standardised over the relevant estimation sample. All specifications include respondent and survey-wave effects. The headline correlated-random-effects (CRE) specification uses the common sample and control set in Table 2: current log household income, subjective status, home ownership and years in the neighbourhood, with both within-person deviations and respondent means entered for every time-varying regressor. The first row and the restriction to observations at or after the first income report are parsimonious models without additional controls; the other Panel A sensitivity rows use the headline control set. "High" noise means "sometimes", "many times" or "always", compared with "never" or "almost never". The categorical row reports a joint test of the four baseline-income interactions with exposure-category indicators. Panel B applies Holm's step-down adjustment jointly to its twelve outcome–moderator tests; the final column reports adjusted p-values in life-satisfaction/sleep/depression order. Panel C evaluates the headline within-person noise slope at one standard deviation below and above mean baseline income; its p-values test each implied slope against zero. Standard errors are clustered by respondent throughout. The separate income and socioeconomic-position decomposition was not prespecified: the prespecified heterogeneity analysis used a composite socioeconomic-status index. The income interactions do not survive the twelve-test Holm adjustment; the life-satisfaction interaction with subjective class does ($p_{\text{Holm}} = 0.016$). Panel B therefore does not support a general material-resource gradient across outcomes.

Appendix B Structured qualitative coding: construction and audit trail

Appendix B.1 Role in the qualitative analysis

The interview study combines the human-led thematic and case-oriented analysis described in Section 4.3 with a structured, episode-level qualitative content analysis. The structured dataset was produced through an *LLM-assisted, human-adjudicated* workflow: a large language model organised candidate evidence under a researcher-defined schema, deterministic programs applied documented checks and repairs, and the study anthropologist and research team retained final authority over coding, case comparison and substantive interpretation.

Organising narrative evidence into categorical fields for systematic comparison is a form of *quantitising* (Sandelowski et al. 2009, Karjus 2025). In this study, quantitising supports rather than replaces the reflexive thematic analysis. The structured dataset is used to describe patterns within the interview corpus, identify contrasting and negative cases, locate transcript evidence and construct session-level indicators for comparison with the panel findings. The themes and theoretical interpretation were developed through human reading, comparison and discussion rather than generated from the structured counts.

The extraction unit is an *episode*: either a bounded noise event or a recurring source–context pattern. An interview can therefore contain several episodes involving different sources, periods, dwellings or appraisals. This preserves contrasts within a case, such as tolerated daytime footsteps and distressing night-time parties. For descriptive comparisons and exploratory tests, however, the interview or household session is the analytic cluster. Episode-level coding retains the multiplicity of experiences within a session; it does not treat separate mentions by the same respondent as independent respondents.

Appendix B.2 Development and final production chain

The workflow was developed iteratively. Scripts 1–4 produced an interview-level prototype, including a developmental model-proposed emotion codebook. Scripts 5–6 introduced early versions of the episode architecture. These stages informed schema development and diagnostic analysis, but the reported structured dataset derives from the later V5.4 production chain. Three programs used in the final chain were stored outside the numbered developmental sequence: `verify_interviewer.py`, `interview_bias_audit.py` and `merge_xls.py`.

Table 9 presents the complete production and provenance chain. It combines the conceptual stages of the workflow with the specific programs, recorded outputs and points of human judgement.

The interviewer-role and elicitation audits used the OpenAI API model identifier `gpt-5.2`. An earlier episode extractor also used `gpt-5.2`; after the model version changed, the corpus was rerun and the final extractor used `gpt-5.4`. The developmental emotion-codebook proposal used `gpt-4.1-mini`. The source code sets these identifiers directly and calls the OpenAI Chat Completions API, so the workflow is reported by API model identifier rather than by ChatGPT subscription tier.

The extractor scripts and output timestamps date the `gpt-5.2` and `gpt-5.4` episode runs to 6 March 2026. Deterministic repair of the final output was completed on 10 March 2026. All

Table 9: Construction and provenance of the structured episode dataset

Stage	Operation	Recorded output or audit	Human role
Schema development	Interview-level extraction, a developmental model-proposed emotion codebook and early episode architectures	Versioned schemas, prompts, corpus-derived examples and diagnostic workbooks	Researchers defined the questions, revised the constructs and selected the episode rather than the interview summary as the final unit
Speaker and elicitation audit	`verify_interviewer.py` classified speaker roles; `interview_bias_audit.py` reviewed potentially leading exchanges	45 high-confidence interviewer-role classifications and 452 elicitation snippets, including eight high-risk flags	Flags were retained for review and incorporated into the later evidence audit rather than treated as automatic exclusions
Episode extraction	`extract_storylinks` mapped `V5.4.py` each transcript to source–pathway–outcome–action episodes under a schema-constrained prompt	Candidate long-form and collapsed episode files, verbatim Spanish evidence with speaker attribution, confidence flags, and invalid-output, dropped-row and weak-mention logs	The model generated candidate evidence structures; researchers retained responsibility for construct definitions and thematic interpretation
Rule-based audit and repair	`post_extraction_ episode_repair.py` normalised labels, applied evidence rules, repaired source and time fields, filtered weak mentions and clustered duplicate slices	222 logged field changes across 141 candidate episodes and four manual-review candidates	Ambiguous cases were surfaced explicitly rather than resolved through unrecorded model judgement
Merge and human adjudication	`merge_xls.py` linked repaired episodes to the pre-interview questionnaire; the anthropologist then reviewed transcript evidence and the final ontology	37 previously unknown pathways adjudicated; 42 conflict labels harmonised; 80 adjudication decisions documented, producing 79 changed workbook cells	The anthropologist held final coding authority for the reviewed pathways and ontology decisions
Reconciliation and versioning	Questionnaire, transcript and episode identifiers were reconciled; the empty questionnaire shell was removed; session-level derivatives and validation logs were produced	Versioned analytic dataset containing 343 substantive episodes in 45 interview or household sessions, with checksums and a session-level analytic file	The research team determined the neighbour-episode definition, publication groupings, case matrix and uses of the structured data

Notes: The prepared plain-text corpus was created from the de-identified transcripts before the programs shown here were run; that document-preparation step is not represented by a supplied conversion script. The final workbook contains no link to the raw audio. Counts describe the recorded workflow and audit files rather than an estimate of overall model accuracy.

model calls used temperature zero. This reduces sampling variation, although a hosted model may still change over time.

Appendix B.3 Episode definition and extraction rules

An episode was defined before the final extraction as either a concrete, bounded event or a recurring source–context pattern. A new episode required a change in source, actor, place, residence, period or lived setting rather than simply the addition of another consequence of the same event. General background vulnerability—for example, being a light sleeper or feeling stressed—was not coded as an episode unless it was connected to a particular sound or recurring

source pattern. Conversely, an audible event described as acceptable, normal or not bothersome remained a valid episode. This rule preserves benign, adverse and mixed appraisals within the same interview.

The final extractor processed the complete interview whenever its length did not exceed 60,000 characters. Longer interviews were divided into 26,000-character segments with 3,000 characters of overlap and a maximum of 20 segments. API failures and invalid JSON responses were retried and logged. Full transcript text was removed from the stored JSON output after extraction.

Every candidate episode was required to contain:

1. a literal trigger category and a short Spanish description;
2. event or recurring-pattern status;
3. the source locus and reported time;
4. one or more pathways and outcomes;
5. an action, response or constraint;
6. a high, medium or low model-confidence flag; and
7. a verbatim Spanish evidence passage with speaker attribution.

Trigger categories distinguished neighbour parties, music, voices, television or audio equipment, footsteps and domestic appliances from construction, building machinery, traffic, sirens, public social activity, animals, nature and violent or explosive sounds. Pathway and outcome fields covered sleep disruption, stress or hypervigilance, fear or insecurity, concentration, social conflict, self-described health symptoms, hopelessness or demoralisation, annoyance or irritation, and benign appraisal. Action fields distinguished direct confrontation, action constrained by fear or social norms, effective and ineffective institutional responses, private adaptation and `no_action_mentioned`.

The code `no_action_mentioned` means that the supporting passage contains no information about action. It is distinct from `unknown`, which denotes a genuinely ambiguous classification. Model confidence is retained only as an audit flag and is not used as a probability or analytic weight. The primary evidence link is the Spanish passage associated with each retained category, allowing the coding to be checked against the respondent's own account.

Appendix B.4 Deterministic audit, human adjudication and versioning

After extraction, deterministic programs normalised the permitted label sets, tightened classifications using explicit textual cues, checked that the evidence contained respondent speech rather than only an interviewer prompt, overlaid elicitation-risk flags, removed unevaluated background mentions, retained explicit benign episodes and combined duplicate segment slices into episode clusters.

The repair program recorded 222 field changes across 141 candidate episodes: 88 changes to outcomes, 45 to the literal source, 40 to time, 40 to pathways, eight to event-versus-recurrence status and one to locus. Four cases were placed on a separate manual-review list.

The study anthropologist then conducted targeted adjudication of every episode whose pathway remained unknown and reviewed the proposed ontology harmonisation. For each of the 37 unknown-pathway episodes, he reread the supporting Spanish evidence. Thirty-six final decisions agreed with the provisional model-assisted recommendation. In case B43_5, a fireworks/construction-noise episode, he replaced benign appraisal with annoyance or irritation on the basis of the transcript.

The anthropologist also approved the harmonisation of 42 occurrences of `interpersonal_conflict` to `social_conflict`. These 37 pathway adjudications and 42 conflict harmonisations produced 79 changed workbook cells. The adjudication log contains an eightieth decision recording the choice to retain `no_action_mentioned` as analytically distinct from `unknown` across 191 episodes. Because that decision preserved the existing values, it produced no additional changed cell. The audit trail therefore records 80 adjudication decisions and 79 data-cell changes.

A cell-by-cell comparison of the pre- and post-adjudication workbooks preserved all 344 candidate rows and 57 fields during this review. Final reconciliation then removed one questionnaire record with no transcript or substantive episode, reconciled the two multi-resident sessions as single analytic clusters, corrected remaining coding and metadata discrepancies, and produced a versioned analytic dataset containing 343 substantive episodes in 45 interview or household sessions.

The final coding scheme retains `fireworks_firecrackers` and `gunshots_bullets` as separate trigger categories. Publication figures combine them as "Fireworks / gunshots (crime signalling)", a grouping for reporting rather than a recoding of the underlying coded dataset. Following final trigger reconciliation, this grouping contains 19 episodes, nine of which—47.4 percent—carry fear or insecurity. This percentage describes the coded episodes in that grouping; the session-level comparisons in the paper continue to use interview or household sessions as the analytic unit.

Appendix B.5 Derived analytic definitions and uses

For comparison with the panel evidence, a *neighbour episode* is defined as an episode attributable to another household or an immediate neighbour. Wider-area traffic and construction, public disorder and other ambient sources are excluded unless the sound is located in a neighbouring household. This definition yields 106 neighbour episodes distributed across 43 of the 45 sessions.

The structured dataset is used in four principal ways. First, it supports within-corpus descriptions of source, timing, recurrence, pathways, outcomes and responses. Second, it makes it possible to preserve multiple and contrasting appraisals within the same interview while returning to the session as the unit for reported percentages. Third, it supports purposive case selection: a 14-session case matrix was constructed for maximum analytic variation and negative-case analysis. Fourth, it provides a systematic bridge to the quantitative evidence by constructing session-level indicators for sleep, stress, social conflict, agency, constraint and related processes.

The closed-form interview questionnaire is used for one additional comparison. Sessions are classified according to whether reported nuisance from neighbour voices, music or footsteps reaches four on a five-point scale. The paper compares the frequency of qualitative sleep and stress evidence across the higher- and lower-scoring groups. These comparisons describe convergence within the purposively recruited interview corpus; they are not used to estimate the prevalence of neighbour noise or its consequences in Temuco or Chile.

Every quotation reported in the paper was rechecked against the de-identified transcript. Only de-identified transcripts were submitted to the API, and potentially identifying details were removed from reported passages.

Appendix B.6 Trustworthiness, privacy and reproducibility

The workflow follows emerging guidance for model-assisted qualitative research by reporting the model and prompt role, preserving an audit trail, linking classifications to source passages, verifying quotations and keeping substantive interpretation under human control (Morgan 2023, Wachinger et al. 2025, Jones 2025, Hill et al. 2026). The deterministic change log, retained evidence passages, targeted anthropologist adjudication and versioned analytic release make each reported coding stage inspectable.

The human adjudication was targeted at unresolved pathways and ontology decisions rather than conducted as a blind second coding of a randomly selected reference set. It therefore documents substantive review and correction but does not estimate an overall model error rate. This is consistent with the supporting role assigned to the structured dataset in the paper. Were the coding workflow itself to become a separate methodological contribution, a useful further step would be blind, independent episode and field coding on a stratified subset of interviews, including examination of whether errors vary across respondent or interview characteristics (Prescott et al. 2024, Ashwin et al. 2026, Hill et al. 2026).

The study received approval from the London School of Economics and Political Science Research Ethics Committee on 25 October 2024, before fieldwork began in December 2024. The transcript corpus was de-identified before API processing, full transcript text was not retained in the extraction output, and quotations were screened for identifying detail.

The planned public replication package can contain the extraction and audit code, prompt templates, empty schemas, the final codebook, aggregate audit counts, checksums and synthetic examples. It will not contain confidential transcripts, raw audio or evidence passages that could permit re-identification.