

Informal Caregiving and the Mental Health of the Young: Panel Evidence from Uruguay

Evelin Lasarga

e1r42@alu.ua.es

University of Alicante

Research Article

Keywords: mental health, youth, caregiving, selection, socioeconomic gradient, Uruguay

Posted Date: September 30th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10952356/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Informal Caregiving and the Mental Health of the Young: Panel Evidence from Uruguay

Evelin Lasarga

University of Alicante — DEcIDE, Interuniversity PhD in Economics, Alicante, Spain

Corresponding author. E-mail: elr42@alu.ua.es

ORCID: 0000-0003-1394-6682

September 9, 2026

Abstract

This paper uses a two-wave panel of Uruguayan youth, followed from ages 12–15 in 2018 to 16–19 in 2022, to map the sociodemographic gradient of mental health across the transition to early adulthood and to ask how much of the association between unpaid caregiving and mental health is causal. The gradient runs along gender and family structure rather than income: young women and youth in father-absent households report substantially worse mental health, while the income gradient familiar from the adult literature is essentially flat. About half of the strong cross-sectional association between caregiving and worse mental health is selection; exploiting within-person variation, a significant effect survives, driven by the *onset* of caregiving—about a third of a standard deviation—concentrated in the care of an older or ill relative and corroborated by a bereavement benchmark. The effect is specific to well-being and does not extend to education, which guards against a selection interpretation. The estimates are treated as credibly identified, quasi-experimental evidence, and their sensitivity to unobserved selection and attrition is bounded. The paper provides the first panel evidence on youth caregiving and mental health for Latin America and a transparent, cross-wave-comparable mental-health measure.

Keywords: mental health; youth; caregiving; selection; socioeconomic gradient; Uruguay.

JEL: I14, I12, J13, J14, I31.

1 Introduction

The mental health of adolescents and young people has become a first-order concern for both public health and economics. Common mental disorders typically emerge before adulthood, are strongly predictive of later education, employment and earnings, and account for a large share of the global burden of disease among the young ([Ridley et al., 2020](#)). Yet economic research on the determinants of youth mental health remains thin, and thinner still for Latin America, where representative longitudinal data have until recently been scarce. This paper uses a new two-wave panel of Uruguayan youth to do two things: to document the sociodemographic gradient of mental health as adolescents move into early adulthood, and to ask how much of the strong observed association between unpaid caregiving and worse mental health reflects a causal effect of taking on care, as opposed to selection.

Uruguay is a setting where these concerns are especially acute. It has the highest suicide rate in Latin America—roughly double the regional average—and, unlike the global pattern in which suicide ranks third or fourth among causes of death in this age range, in Uruguay it is the leading cause of death among adolescents and young adults ([Ministerio de Salud Pública, 2023](#); [World Health Organization, 2024](#)). The adolescent burden is rising: the suicide rate among 15–19 year-olds climbed from about 11 per 100,000 in 2019 to roughly 18 per 100,000 in 2022, and the historically large male–female gap has begun to narrow as young women’s rates rise. Youth mental health in Uruguay is thus not merely an understudied research question but a first-order public-health emergency—one that prompted a national mental-health plan and that lends the developmental window studied here particular urgency.

A large body of work in economics documents robust socioeconomic gradients in health. Income, parental education and early-life conditions shape physical and mental health across the life cycle, and these gaps open early in childhood and widen with age ([Case et al., 2002](#); [Currie, 2009](#); [Cutler et al., 2008](#); [Currie and Vogl, 2013](#); [Almond et al., 2018](#)). For mental health specifically, recent causal evidence links poverty to depression and anxiety through stress and scarcity mechanisms ([Ridley et al., 2020](#)). Two features of this literature motivate the paper’s focus. First, most of the evidence comes from high-income countries and from adults; whether the same gradients hold for adolescents in a middle-income Latin American setting is largely an open empirical question. Second, the gradient need not be driven by income alone: gender, family structure and the demands placed on young people at home may matter as much or more during a developmental window in which most mental-health problems first appear.

This paper studies Uruguay using the ENAJ Youth Panel, which followed a cohort first interviewed in 2018 at ages 12 to 15 and re-interviewed in 2022 at ages 16 to 19. The panel spans a pivotal developmental transition and, unusually for the region, allows the same individuals to be observed before and after they take on adult-like roles. The central research question is: to what extent is the cross-sectional association between unpaid caregiving and adolescent mental health causal, and how does it fit within the broader sociodemographic gradient of youth mental health in Uruguay? The analysis is framed around what individual-

level panel data can credibly identify, and is explicit about what they cannot.

Unpaid caregiving is a natural and policy-relevant lens. A growing share of young people provide informal care, yet they remain largely invisible in official statistics and policy (Miller et al., 2024). Caregiving plausibly affects well-being through time constraints and chronic stress, but it is not randomly assigned: it is allocated within households in response to needs—a new sibling, an aging or ill relative, a parent’s departure—that may themselves bear on a young person’s mental health. The cross-sectional fact that caregivers report worse mental health is therefore difficult to interpret. The panel structure nets out all time-invariant individual and family characteristics by comparing each person to themselves over time, isolating the effect of the transition into caregiving from the standing characteristics of those who care.

The findings can be summarized in four calibrated statements. First, the gradient of youth mental health in Uruguay is steep along gender and family structure but essentially flat along household socioeconomic status: young women report markedly higher distress than young men, paternal presence is protective, and—contrary to the canonical adult income–health gradient—there is no significant association with socioeconomic status (the INSE index as a proxy) once other characteristics are accounted for. Second, about half of the cross-sectional caregiving gap reflects selection; a significant within-person effect survives, driven by the onset of caregiving, which raises a standardized distress index by about a third of a standard deviation and concentrates among those who begin caring for an older or ill relative. A plausibly exogenous bereavement shock moves mental health by a comparable magnitude, corroborating this channel. Third, the effect is specific to well-being: caregiving shows no association with educational outcomes, in contrast to the United States (Miller et al., 2024), which argues against the interpretation that young carers are simply children from disadvantaged or chaotic households. Fourth, apparent gender differences in the caregiving effect are not statistically robust and are better explained by what young people care for than by their sex.

The paper is deliberately cautious about causal language. With two waves it is not possible to test for pre-trends, and within-person estimates remain vulnerable to time-varying shocks that coincide with caregiving onset. The caregiving estimates are therefore treated as credibly identified, quasi-experimental evidence rather than as proof of a causal effect, and subjected to a battery of checks: baseline balance, bounds on selection on unobservables (Oster, 2019), inverse-probability weighting for panel attrition, and tests of robustness to the definition of the outcome. Throughout, the language is matched to what the design can support.

The paper makes three contributions. It provides, to the author’s knowledge, the first panel evidence on caregiving and mental health among youth in Latin America, a population and a region that are both under-studied in the economics of mental health. It offers a transparent, cross-wave-comparable measure of mental health built with item-response methods and validated against the use of mental-health services, addressing a recurrent measurement concern in survey-based work. And it brings discipline to the interpretation of caregiving

correlations by decomposing them into selection and causal components, while documenting the prevalence and sociodemographic profile of young caregivers in Uruguay in the spirit of [Miller et al. \(2024\)](#). The remainder of the paper is organized as follows. Section 2 lays out the conceptual framework; Section 3 reviews the related literature; Section 4 describes the data and the construction of the mental-health measures; Section 5 presents descriptive evidence; Section 6 the empirical strategy; Section 7 the results; and Section 8 concludes.

2 Conceptual Framework

The analysis is organized around three bodies of work: the economics of the socioeconomic gradient in health, the literature on early-life conditions and human capital ([Currie, 2009](#); [Almond et al., 2018](#)), and the care economy, which studies how unpaid care is produced and allocated within households along gendered lines ([Folbre, 2006](#)). The last lens is central here: caregiving is treated not as an individual trait but as household time and effort that fall unevenly on the young, and disproportionately on young women. Adolescence is a sensitive period: it is when most lifetime mental-health conditions first emerge, when human capital is actively accumulated, and when young people begin to take on adult roles. Shocks during this window can therefore have outsized and persistent consequences ([Currie and Vogl, 2013](#)).

Consider a simple framing in which a young person allocates a fixed endowment of time and psychological resources across schooling, leisure, social ties and, increasingly, care for others. Taking on unpaid caregiving imposes two costs: a time constraint that competes with study and rest, and exposure to chronic stress, particularly when the care responds to a relative’s illness or decline. Both channels map directly onto the mechanisms through which adverse circumstances are thought to affect mental health—time scarcity, financial strain and sustained stress ([Ridley et al., 2020](#)). This framing yields a key distinction between the *stock* of being a caregiver and the *flow* of becoming one: if the burden operates through the disruption of taking on a new role, it is the transition—the onset—rather than the standing status that should move mental health, and the effect need not be symmetric when care ends.

The central inferential problem is selection. Caregiving is assigned within households in response to needs that are correlated with unobserved family circumstances—economic precarity, family conflict, the health of other members—each of which may independently affect a young person’s mental health. A cross-sectional comparison of carers and non-carers thus confounds the effect of caregiving with the characteristics of the households that generate it. Comparing each individual to themselves over time removes all time-invariant sources of this selection. What remains is the possibility that a time-varying shock—most plausibly the illness or loss of a relative—drives both the onset of care and the deterioration in mental health. Rather than treating this as a nuisance, the paper uses it: the caregiving estimate is benchmarked against the effect of a discrete bereavement or serious-illness shock to a family member, which the literature treats as a comparatively exogenous source of variation in family circumstances ([Persson and Rossin-Slater, 2018](#); [Böckerman et al., 2023](#)).

This framework also disciplines what is tested. The analysis asks three questions the data can answer: how the sociodemographic gradient of youth mental health looks in Uruguay; how much of the caregiving–mental health association is selection rather than causation; and whether any effect is specific to well-being or extends to educational investment. The last question provides a falsification check: if young carers were simply drawn from disadvantaged households, their education should suffer alongside their mental health; finding an effect on the latter but not the former points to a specific burden of care rather than a generic marker of disadvantage.

3 Related Literature

This paper sits at the intersection of three strands of economic research. The first is the large literature on the socioeconomic gradient in health. [Case et al. \(2002\)](#) show that the income–health gradient is already present in childhood and steepens with age; [Currie \(2009\)](#), [Cutler et al. \(2008\)](#) and [Currie and Vogl \(2013\)](#) document how income, parental education and early-life conditions shape health and human capital across the life cycle, and [Almond et al. \(2018\)](#) survey the long-run consequences of childhood circumstances. This work establishes gradients as a central object of study, but it is overwhelmingly concerned with physical health, with high-income countries, and with adults or young children rather than adolescents.

A more specific body of epidemiological and psychological research documents the gradient in mental health among children and adolescents themselves. Systematic and meta-analytic reviews find that low socioeconomic status is associated with more mental-health problems in youth, but the association is considerably more modest than the adult income–health gradient and is detected more consistently for subjective social status than for objective income ([Reiss, 2013](#); [Quon and McGrath, 2014](#)). A parallel literature, drawing on large cross-national samples, documents a pronounced and widening gender gap: adolescent girls report worse mental health than boys, with the gap larger—paradoxically—in more gender-equal and higher-income countries ([Campbell et al., 2021](#)). This evidence is, however, almost entirely from high-income countries and is overwhelmingly cross-sectional. This paper provides the first individual-level panel evidence on the youth mental-health gradient for a middle-income Latin American setting, and finds the steep gender and family-structure gradients familiar from that literature alongside a flat objective-income gradient—consistent with the weak role of objective socioeconomic status in youth mental health, and in sharp contrast to the canonical adult income–health gradient.

The second strand is the smaller but rapidly growing economics of mental health. [Ridley et al. \(2020\)](#) synthesize evidence on the bidirectional causal relationship between poverty and common mental disorders and the mechanisms—financial strain, stress and scarcity—that link them, and show that anti-poverty interventions such as cash transfers can improve mental health (see also [Baird et al., 2013](#)). Relative to this literature, the paper contributes evidence for a population (youth) and a region (Latin America) that remain markedly understudied.

The third strand exploits discrete family shocks to identify effects on mental health. [Persson and Rossin-Slater \(2018\)](#) use the death of a maternal relative during pregnancy as a plausibly exogenous stressor and find lasting effects on the next generation’s mental health; [Böckerman et al. \(2023\)](#) estimate the mental-health consequences of parental death using register-based event studies. The paper borrows the logic of this design—benchmarking against a comparatively exogenous bereavement or serious-illness shock—to interpret the effect of caregiving onset.

Finally, this paper connects to the nascent economic and demographic literature on young caregivers. [Miller et al. \(2024\)](#) estimate the prevalence of youth and young-adult caregivers in the United States and document their association with reduced educational investment, noting that—unlike adult caregiving—young men and women are nearly equally likely to provide care. Their prevalence and sociodemographic exercises are replicated for Uruguay and extended in two directions: to mental health as an outcome, and to a within-person panel design. Evidence for Latin America is especially scarce; existing Uruguayan work is largely descriptive ([UNICEF-INJU-MIDES, 2023](#); [Giachetto Larraz et al., 2023](#)), and this panel provides the first individual-level longitudinal evidence on caregiving and youth mental health for the country.

A recent and rapidly growing literature, enabled by the inclusion of caregiving questions in large longitudinal surveys, examines the mental-health consequences of young caregiving directly. A systematic review documents small-to-moderate adverse associations that scale with the intensity of care ([Bowman Grangel et al., 2025](#)), and quasi-experimental longitudinal studies in Australian cohorts—using within-person variation ([Fleitas Alfonzo et al., 2024](#)) and augmented inverse-probability weighting ([King et al., 2021](#))—find that taking on informal care worsens adolescent and young-adult mental health. This evidence is, however, confined to high-income countries and typically identifies a dose-response gradient in the hours of care. The present paper extends it in two ways: to a middle-income Latin American setting, where evidence on youth mental health is itself recent ([Bryson and Blanchflower, 2025](#)); and to the extensive margin, locating the effect in the *onset* of caregiving rather than in its intensity—pointing to the disruption of assuming the role rather than the hours it consumes.

A closely related economic literature studies the health consequences of informal caregiving among adults. Using instrumental-variable and panel fixed-effects designs, [Coe and Van Houtven \(2009\)](#) find that caring for an elderly parent raises depressive symptoms, and [Schmitz and Westphal \(2015\)](#) show sizable short-run effects on caregivers’ health that may fade over time; [Bom et al. \(2019\)](#) synthesize this evidence and [Van Houtven et al. \(2013\)](#) document effects on work and wages. This work confronts the same identification problem faced here—caregiving is selected—and motivates the within-person approach, but it has focused almost entirely on adult carers of the elderly. The parallel literature on “young carers” in public health and social policy ([Becker, 2007](#); [Leu and Becker, 2017](#)) establishes that children and adolescents provide substantial family care worldwide, including in developing contexts, yet it is largely descriptive and rarely identifies effects. This paper brings the identification tools of the former to the population of the latter, with mental health as the

outcome; its reading of caregiving as a gendered allocation of household time also connects to the care-economy literature (Folbre, 2006).

4 Data and Measurement

4.1 The ENAJ Youth Panel

This paper uses the Youth Panel of Uruguay’s National Survey on Adolescence and Youth (ENAJ), which followed a cohort first interviewed in 2018 at ages 12 to 15 and re-interviewed in 2022 at ages 16 to 19. The panel was implemented by the National Youth Institute (INJU) with the National Statistics Institute (INE) and UNFPA, and is the first longitudinal study of youth trajectories of its kind in the region. The analysis sample comprises the 959 individuals observed in both waves. Because the 2022 wave was fielded by telephone, with a self-administered module subject to substantial non-response, the telephone instrument is used throughout; all estimates use the panel weights, which adjust for the complex design and for attrition (a person-level longitudinal weight applied to both of an individual’s observations). Attrition is assessed explicitly in Section 6, and its consequences for the estimates are bounded.

4.2 Measuring mental health

The survey does not include a validated psychometric instrument. Its value for present purposes lies instead in a set of items, repeated across both waves, that map onto well-established symptom domains. The measures are built from six such items sharing a common twelve-month reference window: three frequency items on a five-point scale (how often the respondent felt lonely, was so worried as to be unable to sleep, or felt afraid) and three binary items (whether, in the last twelve months, the respondent felt sad or hopeless for two weeks or more, self-harmed intentionally, or seriously considered suicide). Every item is coded so that higher values denote worse mental health, eliminating the direction inconsistencies that arise when only some items are reverse-scored.

Comparability across waves is essential because identification rests on within-person change. The items share the same response format and a common twelve-month reference window in 2018 and 2022; rather than assume they measure distress identically over time, their cross-wave comparability is assessed directly through the measurement-invariance tests reported below. To summarize the items into a single index, one item-response model is estimated on the pooled two-wave data, treating the five-point items as graded-response items and the binary items as two-parameter items. Estimating a single model imposes common item parameters across waves, which places the 2018 and 2022 latent scores on the same metric by construction. Expected-a-posteriori (EAP) scores are extracted and standardized to have unit standard deviation, so that all effects are reported in standard-deviation units.

This construction is preferable to two alternatives that are deliberately avoided. Estimat-

ing a factor or item-response model separately within each wave and rescaling the scores to a zero-to-one-hundred range—a common practice—produces scores that are not comparable across waves and whose magnitudes are sensitive to a single extreme observation, rendering changes over time difficult to interpret. An unweighted sum of the items, by contrast, treats all items as equally informative; the item-response model instead weights each item by its discrimination, which the data show varies substantially (Figure 1): the fear item discriminates weakly while the suicidal-ideation item discriminates strongly.

The index has three documented properties. First, on dimensionality, a parallel analysis of the polychoric correlations supports a single dominant factor, with a weaker second facet distinguishing internalizing symptoms from self-harm and suicidality; the general factor is used and a two-facet specification is reported as a robustness check. Second, on measurement invariance, the equal-loadings (metric) model is not rejected against the configural one ($p = 0.09$), supporting a common cross-wave metric; full scalar (equal-threshold) invariance is rejected, as is common in practice, so the analysis relies on the anchored, metric-invariant scores. Third, on reliability, ordinal omega is approximately 0.93 and the model-based (IRT) marginal reliability is 0.71.

The index is validated against external markers of mental-health need that are not used in its construction (Table 1). The index is positively correlated with having seen a psychologist or psychiatrist, with the use of psychotropic medication, and with a professional mental-health diagnosis (correlations between roughly 0.36 and 0.45); the probability of a professional diagnosis is several times higher among those classified as severely distressed than among the rest (correlations of roughly 0.30 to 0.38). Reassuringly, the professional diagnosis—the marker least mechanically tied to the self-reported symptoms and the most demanding—shows the strongest association. The index therefore tracks clinically relevant help-seeking rather than noise. For analyses that require a discrete outcome, a transparent indicator of severe distress is defined (self-harm, suicidal ideation, or a two-week episode of sadness or hopelessness), in preference to a multidimensional poverty threshold whose lenient cut-off would mechanically inflate prevalence; the central results are shown to be robust to using, instead, a transparently constructed multidimensional mental-health-poverty index in the spirit of Bourguignon and Chakravarty (2003).




Figure 1: Estimated item discriminations from the cross-wave item-response model. The index weights items by discrimination (fear weak, suicidal ideation strong), unlike an equal-weight sum.

4.3 Caregiving and other variables

Both waves include a care module asking whether the respondent provided unpaid care in the previous twelve months, and, for caregivers, the relationship to the care recipient, the recipient’s age and disability status, and weekly hours of care. A caregiving indicator is defined from these items and, exploiting the panel, the *onset* of caregiving—a transition from non-caregiver to caregiver between waves—together with its cessation. The relationship and disability information distinguish care for an older or ill relative from care for a younger sibling or one’s own child. These are complemented with household composition; the INSE household socioeconomic index, which is available and on a common scale in both waves (the 2018 and 2022 indices correlate 0.73) and which is standardized within wave; standard demographics; and—measured retrospectively in 2022—a set of pandemic-era family shocks, including the serious illness or death of a relative, used as a benchmark exogenous shock.

Table 1: The mental-health index: discrimination, reliability and external validity

Item / statistic	Value
IRT discrimination (α): loneliness / sleep / fear	1.89 / 1.80 / 0.99
IRT discrimination (α): sad-2wk / self-harm / ideation	2.22 / 2.32 / 3.39
Ordinal omega / IRT marginal reliability	0.93 / 0.71
Corr. with psychiatrist / medication / diagnosis	0.30 / 0.32 / 0.38
$P(\text{diagnosis} \mid \text{severe})$ vs $P(\text{diagnosis} \mid \text{not severe})$	0.21 vs 0.03

5 Descriptive Evidence

Before turning to identification, this section describes the cohort, the evolution of mental health across the transition to early adulthood, the sociodemographic gradient, and the place of caregiving in young people’s lives. Three facts motivate the analysis that follows: mental health deteriorates and polarizes as the cohort ages; the gradient runs along gender and family structure rather than income; and caregivers differ markedly from non-caregivers in ways that the raw data cannot attribute to caregiving itself.

5.1 The cohort and the deterioration of mental health

The panel follows 959 young people from ages 12–15 in 2018 to 16–19 in 2022; just under half are women and roughly one in five identify as afro-descendant (Table 2). Over these four years their mental health worsens on every summary measure. The standardized distress index rises by about a third of a standard deviation at the mean, but the mean understates the change: the distribution shifts and fans out (Figure 2). Its standard deviation grows from 0.92 to 1.05 and the gap between the ninetieth and tenth percentiles widens from 2.37 to 3.08, with essentially all of the movement in the upper (worse) tail—the tenth percentile

is unchanged while the ninetieth rises sharply. The share of young people in severe distress almost triples, from 9.8 to 28.2 percent. The item-level pattern is consistent: loneliness and sleep problems become more frequent, and the binary markers of serious distress rise sharply (sadness or hopelessness lasting two weeks from 8.6 to 23.4 percent, self-harm from 2.8 to 7.6 percent, and suicidal ideation from 2.7 to 12.0 percent), while fear, more tied to the school environment, declines slightly. Use of mental-health services rises in parallel, corroborating that the index tracks real need rather than changes in reporting alone. Yet need substantially outstrips treatment: among the severely distressed in 2022, 55 percent had received any mental-health service—a psychologist or psychiatrist visit, medication, or a professional diagnosis—leaving a treatment gap of about 45 percent, against 18 percent among the rest; and access rises steeply with socioeconomic status (48 percent of severely distressed youth in the lowest INSE tercile received any service, versus 66 percent in the highest). This mirrors, in a middle-income setting, the under-treatment of child mental-health problems documented for the United States ([Horwitz et al., 2012](#)).

Table 2: Sample composition and mental health, by wave (panel, $n = 959$)

	2018	2022
Female (%)	49.2	49.2
Afro-descendant (%)	18.7	18.7
Mean age (years)	13.6	17.5
Distress index, mean (SD units)	−0.21	+0.16
Distress index, std. dev.	0.92	1.05
Distress p10 / p50 / p90	−1.43 / −0.20 / 0.94	−1.43 / +0.09 / 1.65
Severe distress (%)	9.8	28.2
Felt lonely (mean, 1–5)	1.84	2.10
Could not sleep from worry (1–5)	1.82	2.08
Felt afraid (1–5)	2.05	1.89
Sad/hopeless ≥ 2 weeks (%)	8.6	23.4
Self-harm (%)	2.8	7.6
Suicidal ideation (%)	2.7	12.0
Saw psychiatrist (%)	6.5	8.9
Psychotropic medication (%)	6.6	11.2
Professional diagnosis (%)	—	8.1

5.2 The sociodemographic gradient

The deterioration is uneven across groups, and the pattern is the central descriptive finding of the paper (Table 3, Figure 3). The largest and fastest-widening gap is by sex: young men and women start about half a standard deviation apart in 2018 and diverge to about two-thirds of a standard deviation by 2022, as young women’s mental health worsens far more. Family structure matters too—youth in father-absent households fare worse—as does afro-descendant status, though more modestly. By contrast, there is essentially no gradient




Figure 2: The distribution of the distress index shifts rightward and fans out between 2018 and 2022, with the deterioration concentrated in the upper (worse) tail.

by household socioeconomic status: mean distress is flat across socioeconomic terciles in both waves. This stands in sharp contrast to the well-documented income–health gradient among adults and frames the regression evidence in Section 7.

Table 3: Mean distress (SD units) by group and wave

Group	2018	2022
Men	−0.46	−0.15
Women	+0.07	+0.52
Father present (2018)	−0.24	+0.10
Father absent (2018)	−0.15	+0.28
Non-afro	−0.24	+0.14
Afro-descendant	−0.09	+0.25
SES (INSE) — low tercile	−0.14	+0.10
SES (INSE) — middle tercile	−0.14	+0.24
SES (INSE) — high tercile	−0.22	+0.17

5.3 Who provides care, and the raw caregiving gap

Caregiving is a common feature of these young lives. Under the broad definition, 28.8 percent provided unpaid care in 2018 and 15.2 percent in 2022. The decline is concentrated in the care of young children: care of recipients aged 0–12 falls sharply between waves, while care for an older or ill relative is essentially stable (4.3 to 3.9 percent). The fall thus reflects the cohort aging out of the care of young children as it moves into late adolescence, rather than the loss of elderly recipients. The burden is far from token—among caregivers in 2022 the median is twelve hours of care per week, and 45 percent provide twenty hours or more (Table 4). Caregiving is modestly more common in lower-socioeconomic households. Participation in care is near-equal by sex—young women are only slightly more likely than young men to provide any care—but its intensity is not: young women are substantially more likely to




Figure 3: Mean distress by group and wave. The gap widens sharply by sex and, more modestly, by family structure, while it remains essentially flat across socioeconomic (INSE) terciles.

provide intensive care of ten or more hours per week, in both waves (Table 5). Crucially for what follows, caregiving is not a stable state: between waves, 83 young people take on care and 211 relinquish it, while 58 care in both years and 607 in neither. This within-person movement is what the identification strategy exploits.

The raw association between caregiving and mental health is large and grows over time: caregivers score 0.35 standard deviations worse than non-caregivers in 2018 and 0.40 worse in 2022. Taken at face value this looks like clear evidence of harm. But caregivers and non-caregivers differ in many ways—family circumstances, the presence of an ill relative, economic precarity—that bear independently on mental health, and the cross-section cannot separate the effect of caregiving from the characteristics of those who provide it. Disentangling the two is the task of the empirical strategy that follows.

Table 4: Caregiving prevalence, profile and transitions

	Value
Any unpaid care, 2018 / 2022 (%)	28.8 / 15.2
Prevalence by SES tercile, 2018 (low/mid/high, %)	32.2 / 29.5 / 22.5
Hours/week among caregivers, 2022 (median; $\geq 20h$)	12; 45%
Transitions: never / onset / cessation / always	607 / 83 / 211 / 58
Raw distress gap, carers – non-carers (2018 / 2022, SD)	+0.35 / +0.40

6 Empirical Strategy

This paper pursues two complementary objectives: to describe the sociodemographic gradient of youth mental health, and to separate the causal effect of caregiving from selection. The text is explicit about which estimates are descriptive and which are interpreted causally, and calibrates its language to the design throughout.

Table 5: Prevalence of young caregivers, by definition and by age of care recipient

Caregiver definition	2018 (%)	2022 (%)
Any unpaid care, last 12 months	28.8	15.2
Intensive care (≥ 10 h/week)	9.4	10.0
<i>By age of care recipient</i>		
65 years or older	4.3	3.9
13–64 years	1.5	2.8
4–12 years	17.3	7.6
0–3 years	11.5	2.8

Notes: Weighted. Age bands of the care recipient are not mutually exclusive (a respondent may care for recipients in more than one band). The decline in caregiving between waves is concentrated in the care of young children (ages 0–12), while care for an older relative is stable. Participation in any care is near-equal by sex (men 14.1% vs women 16.4% in 2022), but intensive care (≥ 10 h/week) is disproportionately female (men 8.2% vs women 12.0% in 2022, and 6.4% vs 12.7% in 2018). Afro-descendant youth are over-represented among caregivers (26.2% vs 17.4%).

6.1 The gradient

To characterize the gradient, weighted pooled regressions of the standardized distress index on socioeconomic status, sex, family structure and age are estimated, with the survey year included to absorb common trends. Coefficients are reported in standard-deviation units, and, for discrete outcomes, average marginal effects are reported alongside odds ratios. These regressions are descriptive: they summarize how mental health varies across groups, not the effect of manipulating any single characteristic.

6.2 Selection versus causation

The central exercise decomposes the cross-sectional caregiving–mental health association into a part attributable to selection and a part consistent with a causal effect. The pooled cross-sectional association is estimated first; within-person variation is then exploited through a first-difference specification, which with two periods is equivalent to a two-way fixed-effects model and removes all time-invariant individual and household characteristics. Because the effect of taking on care need not mirror the effect of relinquishing it, the onset and the cessation of caregiving are allowed to enter separately:

$$\Delta y_i = \alpha + \beta_{\text{on}} \text{Onset}_i + \beta_{\text{off}} \text{Cease}_i + \Delta X_i' \gamma + \Delta \varepsilon_i, \quad (1)$$

where Δy_i is the within-person change in distress and Onset_i (Cease_i) indicates beginning (ending) caregiving between waves. The parameter of interest, β_{on} , is identified under strict exogeneity: conditional on the time-invariant fixed effect, the timing of onset is assumed uncorrelated with the idiosyncratic evolution of mental health. Because β_{on} is identified only off the individuals who change caregiving status between waves, it is complemented

with a value-added specification estimated on the full sample:

$$y_{i,2022} = \alpha + \lambda y_{i,2018} + \theta \text{Onset}_i + \psi \text{Cease}_i + X_i' \gamma + u_i, \quad (2)$$

which conditions on the baseline level of distress $y_{i,2018}$ rather than differencing it out. The two designs rest on opposite assumptions about how selection operates: first-differencing is consistent when the confounding works through a time-invariant individual trait, whereas the value-added model is consistent when it works through the baseline level of the outcome. When neither holds exactly, the true effect is bracketed by the two estimates (Angrist and Pischke, 2009). A cluster bootstrap of the onset estimate is also reported, given the modest number of switchers.

This assumption is not innocuous, and its limitations should be stated plainly. With only two waves it is not possible to test for parallel pre-trends, and first-differencing does not purge time-varying shocks that coincide with caregiving onset, nor does it by itself rule out reverse causality or anticipation. These threats are addressed in three ways. Selection on observables is assessed by testing whether those who begin caring differ at baseline, particularly in baseline mental health. Selection on unobservables is bounded using the coefficient-stability approach of Oster (2019). And the most plausible time-varying confounder is confronted directly by benchmarking the caregiving estimate against the effect of a discrete bereavement or serious-illness shock to a family member—a comparatively exogenous source of variation (Persson and Rossin-Slater, 2018; Böckerman et al., 2023).

6.3 Robustness, attrition and inference

Robustness is probed along several dimensions. Panel attrition is corrected by the longitudinal survey weights, which INE constructs for the balanced panel precisely to adjust for it; the onset estimate is verified to be robust to this weighting. The onset effect is re-estimated after entropy-balancing those who begin caring against those who do not on baseline covariates. The intensive margin (hours of care) is examined separately from the extensive margin (becoming a caregiver). And q -values are reported that adjust for multiple hypothesis testing across the pre-specified family of within-person causal estimates: caregiving onset, its two sex-specific counterparts, the symmetric first difference, and the bereavement benchmark. Inference uses heteroskedasticity-robust standard errors for the within-person specifications and standard errors clustered at the departmental level where area-level regressors enter.

6.4 What the design does not support

This paper does not pursue an instrumental-variables strategy. The natural instruments for caregiving—the arrival of a dependent in the household, the presence of an older or ill relative—violate the exclusion restriction, because such events plausibly affect a young person’s mental health through channels other than their own caregiving; the corresponding first stage is also weak, and the sample is small. A transparent within-person design with

explicit bounds is therefore preferred to an instrument that could not be defended. For the same reasons, departmental income inequality enters only as a contextual control rather than as an object of inference; it is statistically insignificant and does not affect the estimates.

6.5 Statistical power

Given the sample size, minimum detectable effects are reported so that null results can be interpreted correctly. The standardized continuous index is the primary outcome and is adequately powered to detect effects of moderate size; analyses of discrete outcomes and of subgroups are less well powered and are presented as exploratory, with the convention that an imprecise null reflects absence of evidence rather than evidence of absence.

7 Results

7.1 The sociodemographic gradient

In a pooled regression of distress on socioeconomic status, sex, family structure and age (Table 6), the dominant gradient is by sex: young women report 0.59 standard deviations more distress than young men. Family structure matters too—paternal presence lowers distress by about 0.13 standard deviations—and distress rises modestly with age. By contrast, the association with household socioeconomic status is small and statistically indistinguishable from zero. This flat gradient is not an artifact of measurement or restricted range: it is small and non-significant whether the measure is the INSE, (log) household income, or a categorical socioeconomic level, and the INSE spans its full range in the sample.¹ The mental-health gradient among Uruguayan youth thus runs along gender and family structure rather than income, a contrast with the canonical adult income–health gradient that is itself a substantive finding.

7.2 Selection versus causation

The cross-sectional association between caregiving and mental health is sizable: caregivers score about 0.33 standard deviations worse on the distress index, with or without controls (Table 7). A large part of this gap reflects who becomes a caregiver rather than the effect of caregiving. When each individual is compared to themselves over time, the association falls by about half, to 0.17 standard deviations, but remains statistically significant ($p = 0.03$): roughly half the cross-sectional gap is selection, and a genuine within-person effect survives. That effect is driven entirely by the transition into care: beginning to provide care raises distress by 0.31 standard deviations, while ceasing to provide care has no significant offsetting

¹The flat gradient holds in each wave separately, not only in the pooled specification. In 2018 alone, the association between distress and socioeconomic status is statistically insignificant for both the INSE ($p = 0.17$) and (log) per-capita ECH household income ($p = 0.78$), and remains so once sex, afro-descendant status and age are controlled for.

Table 6: The sociodemographic gradient (outcome: standardized distress index)

Variable	Coef. (SD)	Std. error
Socioeconomic status (INSE, z)	0.035	0.027
Female	0.590***	0.053
Afro-descendant	0.128	0.074
Father present (2018)	−0.126**	0.054
Caregiving	0.345***	0.061
Age	0.091***	0.024
Year 2022	0.083	0.119

Notes: $n = 1,918$ person-years; weighted; robust standard errors. *** $p < 0.01$, ** $p < 0.05$. Descriptive specification.

effect, a striking asymmetry that points to the disruption of taking on the role rather than to caregiving status per se. Figure 4 summarizes the decomposition.

Table 7: Selection vs. causation (outcome: within-person change in distress)

Specification	Coef. (SD)	p -value
Cross-sectional association (with controls)	0.336	<0.01
Within-person first difference (symmetric)	0.169	0.03
Onset of caregiving	0.305	0.01
Cessation of caregiving	−0.102	0.35
Onset, care for older/ill relative	0.545	0.02
Onset, care for sibling	0.331	0.17
Bereavement / illness shock (benchmark)	0.236	0.04
Value-added (LDV), full sample (onset)	0.373	<0.01

Notes: $n = 959$; 83 begin and 211 stop caring between waves. Weighted; robust standard errors. Coefficients in SD units. The value-added (LDV) row regresses 2022 distress on caregiving onset conditional on 2018 distress and covariates (full sample); the onset first-difference estimate has a cluster-bootstrap 95% CI of roughly $[0.05, 0.53]$.

The effect is concentrated in the care of an older or ill relative. Splitting onset by the type of recipient, the deterioration is largest and statistically significant for those who begin caring for an older or ill relative (0.55 standard deviations, $p = 0.02$), not significant for those caring for a sibling, and absent for the small number caring for their own children. This pattern coheres with the benchmark: a retrospectively reported bereavement or serious-illness shock to a family member raises distress by 0.24 standard deviations ($p = 0.04$), a comparably exogenous source of variation that operates through the same channel of family decline.

A natural concern is that the onset effect merely reflects the family-health shock that triggers caregiving—the very channel the benchmark captures—rather than caregiving itself. Three pieces of evidence separate the two. First, most caregiving onset is unrelated to such a shock: only 28 percent of those who begin caring report a serious illness or death in the




Figure 4: The cross-sectional caregiving association shrinks to insignificance under within-person differencing; what survives is the onset effect, concentrated in care for an older or ill relative and mirrored by the bereavement benchmark.

family. Second, the onset effect is essentially unchanged when the bereavement shock is controlled for (0.29, $p = 0.02$), which enters as a separate, additive channel (0.23), and there is no interaction between the two—caregiving matters equally whether or not a family-health shock occurred. Third, and most tellingly, the onset effect remains significant among the majority who report no such shock (0.28, $p = 0.04$). Caregiving and family-health shocks are thus distinct, additive channels; the benchmark corroborates the mechanism rather than confounding it.

Apparent gender differences do not survive scrutiny. Although the onset effect is statistically significant for young men and not for young women when each is estimated separately, the difference between them is not itself significant (the interaction carries a p -value of 0.37), and once the type of care recipient is controlled for the male coefficient collapses to near zero. The seeming concentration among young men reflects that they more often provide the more demanding kinds of care, not a difference by sex; consistent with this, young women who begin caring already report elevated distress at baseline. The caregiving effect is therefore presented as broadly shared, and no gender-specific mechanism is built on it. Finally, the onset effect is not an artifact of the small number of switchers: a value-added model that conditions on baseline distress and uses the full sample of 959 individuals yields an even larger onset coefficient (0.37, $p < 0.01$; Table 7), and a cluster bootstrap of the first-difference estimate excludes zero (95% confidence interval roughly $[0.05, 0.53]$), although its magnitude is imprecisely estimated.

7.3 Robustness

Several checks support a causal reading of the onset effect while delimiting its strength (Table 8); they are grouped here by the threat each addresses.

Selection. Those who begin caring do not differ significantly from others in baseline

mental health (a difference of about 0.11 standard deviations, $p = 0.41$), which is reassuring against selection on the level of the outcome, though they are somewhat more likely to come from father-absent households, which the specification controls for. Reverse causality is likewise unsupported: among those not caring in 2018, baseline distress does not significantly predict beginning to care by 2022, so it is not the case that youth whose mental health is deteriorating are the ones selected into caregiving. Finally, the coefficient-stability diagnostic of [Oster \(2019\)](#) implies that selection on unobservables would need to be several times stronger than selection on observables to overturn the result under a conventional bound.

Attrition. Panel attrition is corrected by the longitudinal survey weights, and the onset estimate is robust to this weighting: it is 0.39 without the weights and 0.31 with them (both $p < 0.05$), so the adjustment moves the estimate little and never toward zero. Entropy-balancing weights that equate starters and non-starters on baseline covariates give a similar figure (0.381).

The pandemic window. A concern specific to this window is that 2018–2022 spans the COVID-19 pandemic, so the deterioration, the family shocks and caregiving all fall in the pandemic period. Two features address this. First, the within-person design differences out any component of the pandemic common to the cohort: the onset estimate compares those who begin caring to those who do not, both of whom lived through the pandemic, so a uniform pandemic effect is absorbed by the survey-year term. Second, controlling for the general pandemic disruptions that could differentially affect caregivers—household income loss, job loss, and remote schooling—leaves the onset estimate essentially unchanged (0.29, $p = 0.03$). The specification does not additionally control for pandemic-era family ruptures (separation, bereavement), which are co-triggers of caregiving rather than confounders and are examined separately above. It should nonetheless be cautioned that the rise in distress *levels* partly reflects the pandemic, and that the retrospective pandemic-shock reports may carry recall bias correlated with current mental state.

Outcome definition and inference. The effect operates on the extensive margin: among caregivers, weekly hours of care are unrelated to distress. The onset estimate survives an adjustment for multiple hypothesis testing ($q = 0.03$) and reproduces when the index is replaced with a transparently constructed multidimensional mental-health-poverty index, confirming that the conclusion does not depend on the choice of outcome. Departmental income inequality, included only as a context control, is insignificant.

7.4 Specificity: caregiving and education

A final result sharpens the interpretation. If young caregivers were simply children from disadvantaged or disorganized households, the burden should be visible in their schooling as well as their mental health. It is not: caregiving is unrelated to dropping out, to grade repetition, or to attained schooling, both cross-sectionally—where these regressions are well powered—and within individuals. The contrast with [Miller et al. \(2024\)](#), who find a robust caregiving–education association in the United States, is instructive: in this setting the burden of care falls on well-being specifically and does not appear as a generic marker of

Table 8: Robustness and specificity

Check	Result
Baseline distress, starters vs others	-0.11 vs -0.22 ($p = 0.41$)
Oster (2019) δ^* to drive onset to zero ($R_{\max} = 1.3R^2$)	≈ 4.8
Onset, unweighted / weighted (attrition)	0.39 / 0.31
Onset with entropy-balancing weights	0.381
Onset q -value (multiple testing, BH)	0.033
Replicates on BCa poverty index	yes
Departmental Gini (context control)	not significant
Caregiving $\rightarrow$ education (dropout/repetition)	null (well powered)

Notes: Minimum detectable effects (80% power): distress 0.34 SD, severe 0.19, dropout 0.08, repetition 0.12.

disadvantage. This specificity is itself evidence against the selection interpretation of the mental-health result.

8 Conclusions

This paper has used a two-wave panel of Uruguayan youth to map the sociodemographic gradient of mental health across the transition to early adulthood and to ask how much of the link between unpaid caregiving and mental health is causal. Three conclusions stand out. First, the gradient of youth mental health in Uruguay runs along gender and family structure rather than income: young women and youth in father-absent households fare markedly worse, while the income gradient so prominent for adults is essentially flat. Second, about half of the strong cross-sectional association between caregiving and worse mental health is selection; what is left, once young people are compared to themselves, is a significant within-person effect, driven by the onset of caregiving, of about a third of a standard deviation, concentrated in the care of an older or ill relative and corroborated by a bereavement benchmark. Third, this burden is specific to well-being and does not extend to education, which both interprets the effect and guards against a selection reading.

This paper has been deliberate about the strength of these claims. Two waves do not permit a test of parallel trends, the outcome is self-reported, and the pandemic-era shocks are measured retrospectively; subgroup analyses are underpowered. The caregiving estimate is therefore presented as credibly identified, quasi-experimental evidence rather than as a definitive causal effect, supported by baseline balance, bounds on unobserved selection, attrition corrections, and robustness to the definition of the outcome.

The findings nonetheless carry clear implications. Young caregivers are numerous and largely invisible to policy, and the mental-health cost of care concentrates among those supporting an ailing relative—precisely the situations a long-term-care system is meant to relieve. Strengthening Uruguay’s national care system so that adolescents are not the residual providers of care for sick and older household members, and integrating mental-health

support into the settings where young people can be reached—educational centers and the digital platforms of the Ceibal program—are responses that follow directly from the evidence. Because the effect is tied to taking on care rather than to caregiving status—it is the onset, not the ongoing state or the number of hours, that moves mental health—support targeted at young people newly entering a caregiving role may be especially valuable. The treatment gap documented here—close to half of severely distressed youth receive no mental-health service, and fewer among the poor—reinforces the case for expanding and equalizing access.

Several extensions would sharpen the analysis. Additional waves of the ENAJ panel would permit tests of pre-trends and of persistence, moving the caregiving estimate closer to a fully causal interpretation; linkage to administrative records on health-service use and education would reduce reliance on self-reports; and a deeper treatment of the care of older and ill relatives—its timing, intensity and resolution—would illuminate the mechanism identified here. More broadly, these results suggest that the economics of youth mental health in Latin America is a fertile and underexplored area, in which even modest longitudinal data can deliver credible and policy-relevant evidence.

Acknowledgements

I thank my supervisor, Josep Peris, for his guidance and comments, and my colleague Santiago Budría for helpful discussions. Part of this research was carried out during a research visit at the University of Rome Tor Vergata, whose hospitality is gratefully acknowledged. I also gratefully acknowledge the 2023 Research Prize of the Instituto Juan Pablo Terra. All remaining errors are my own.

Statements and Declarations

Funding. Financial support from the Spanish Ministerio de Ciencia, Innovación y Universidades (PID2022-138111NB-I00) and the Juan Pablo Terra Institute’s 2023 Research Prize is gratefully acknowledged.

Competing interests. The author declares that she has no financial or non-financial competing interests.

Data availability. The datasets analysed during the current study are openly available in the National Statistics Institute (INE) microdata catalogue (ANDA) at <https://www4.inecub.uy/Anda5/index.php/catalog/739> — the National Survey on Adolescence and Youth (ENAJ) Youth Panel, 2018 and 2022 waves (INJU–MIDES / INE / UNFPA). Replication code is available from the author on reasonable request.

Author contributions. Evelin Lasarga is the sole author and is responsible for all aspects of the study.

Ethics approval. The analysis uses anonymized secondary microdata from an official public survey and did not require additional ethics approval.

A The multidimensional mental-health-poverty (BCa) index

As a robustness check on the outcome definition, a multidimensional mental-health-poverty index is reconstructed transparently, in the tradition of [Bourguignon and Chakravarty \(2003\)](#). Each of the six symptom items defines a deprivation (a person is deprived in the frequency items if they report the symptom “often” or more, and in the binary items if they report it at all). Following the adjusted headcount logic, an individual’s index is the product of the incidence of deprivation and its intensity across dimensions, $BCa_i = \varphi_i(G) \cdot A_i$, where $\varphi_i(G)$ aggregates the normalized deprivation gaps and A_i rewards achievements in non-deprived dimensions; the population index is the mean of BCa_i . Unlike the version in the original submission, the deprivation thresholds and the aggregation are stated explicitly and prevalence is not inflated by counting mild, occasional symptoms as deprivations. The central results—the collapse of the cross-sectional caregiving association under first-differencing and the surviving onset effect—reproduce when this index (continuous, or a binary “poverty” indicator) replaces the item-response score, confirming that the conclusions do not depend on the choice of outcome.

References

- Douglas Almond, Janet Currie, and Valentina Duque. Childhood circumstances and adult outcomes: Act II. *Journal of Economic Literature*, 56(4):1360–1446, 2018.
- Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press, 2009.
- Sarah Baird, Jacobus de Hoop, and Berk Özler. Income shocks and adolescent mental health. *Journal of Human Resources*, 48(2):370–403, 2013.
- Saul Becker. Global perspectives on children’s unpaid caregiving in the family: Research and policy on ‘young carers’ in the uk, australia, the usa and sub-saharan africa. *Global Social Policy*, 7(1):23–50, 2007.
- Petri Böckerman, Mika Haapanen, and Christopher Jepsen. Dark passage: Mental health consequences of parental death. *American Journal of Health Economics*, 9(4):584–604, 2023.
- Judith Bom, Pieter Bakx, Frederik Schut, and Eddy van Doorslaer. The impact of informal caregiving for older adults on the health of various types of caregivers: A systematic review. *The Gerontologist*, 59(5):e629–e642, 2019.
- François Bourguignon and Satya R. Chakravarty. The measurement of multidimensional poverty. *Journal of Economic Inequality*, 1(1):25–49, 2003.
- Aoife Bowman Grangel, Jennifer McMahon, Nikki Dunne, Fay O’Donoghue, and Stephen Gallagher. Health outcomes and psychosocial determinants in young carers: A systematic review. *The Lancet Public Health*, 10(7):e625–e639, 2025.
- Alex Bryson and David G. Blanchflower. The mental health of the young in latin america. *Social Indicators Research*, 2025. Advance online publication.
- Olivia L. K. Campbell, David Bann, and Praveetha Patalay. The gender gap in adolescent mental health: A cross-national investigation of 566,829 adolescents across 73 countries. *SSM - Population Health*, 13:100742, 2021.
- Anne Case, Darren Lubotsky, and Christina Paxson. Economic status and health in childhood: The origins of the gradient. *American Economic Review*, 92(5):1308–1334, 2002.
- Norma B. Coe and Courtney Harold Van Houtven. Caring for mom and neglecting yourself? the health effects of caring for an elderly parent. *Health Economics*, 18(9):991–1010, 2009.
- Janet Currie. Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development. *Journal of Economic Literature*, 47(1):87–122, 2009.

- Janet Currie and Tom Vogl. Early-life health and adult circumstance in developing countries. *Annual Review of Economics*, 5:1–36, 2013.
- David M. Cutler, Adriana Lleras-Muney, and Tom Vogl. Socioeconomic status and health: Dimensions and mechanisms. NBER Working Paper 14333, National Bureau of Economic Research, 2008.
- Ludmila Fleitas Alfonzo, George Disney, Ankur Singh, Koen Simons, and Tania King. The effect of informal caring on mental health among adolescents and young adults in australia: A population-based longitudinal study. *The Lancet Public Health*, 9(1):e26–e34, 2024.
- Nancy Folbre. Measuring care: Gender, empowerment, and the care economy. *Journal of Human Development*, 7(2):183–199, 2006.
- Gustavo Giachetto Larraz et al. Trastornos mentales en la infancia en Uruguay. *Working paper*, 2023.
- Sarah McCue Horwitz, Michael S. Hurlburt, Amy Heneghan, et al. Mental health problems in young children investigated by US child welfare agencies. *Journal of the American Academy of Child & Adolescent Psychiatry*, 51(6):572–581, 2012.
- Tania King, Ankur Singh, and George Disney. Associations between young informal caring and mental health: A prospective observational study using augmented inverse probability weighting. *The Lancet Regional Health – Western Pacific*, 15:100257, 2021.
- Agnes Leu and Saul Becker. A cross-national and comparative classification of in-country awareness and policy responses to ‘young carers’. *Journal of Youth Studies*, 20(6):750–762, 2017.
- Katherine E. M. Miller, Joanna L. Hart, Mateo Useche Rosania, and Norma B. Coe. Youth caregivers of adults in the united states: Prevalence and the association between caregiving and education. *Demography*, 61(3):829–847, 2024.
- Ministerio de Salud Pública. El suicidio adolescente en Uruguay: un asunto de salud pública. Technical report, Ministerio de Salud Pública del Uruguay, Montevideo, 2023.
- Emily Oster. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204, 2019.
- Petra Persson and Maya Rossin-Slater. Family ruptures, stress, and the mental health of the next generation. *American Economic Review*, 108(4–5):1214–1252, 2018.
- Elizabeth C. Quon and Jennifer J. McGrath. Subjective socioeconomic status and adolescent health: A meta-analysis. *Health Psychology*, 33(5):433–447, 2014.
- Franziska Reiss. Socioeconomic inequalities and mental health problems in children and adolescents: A systematic review. *Social Science & Medicine*, 90:24–31, 2013.

- Matthew Ridley, Gautam Rao, Frank Schilbach, and Vikram Patel. Poverty, depression, and anxiety: Causal evidence and mechanisms. *Science*, 370(6522):eaay0214, 2020.
- Hendrik Schmitz and Matthias Westphal. Short- and medium-term effects of informal care provision on female caregivers’ health. *Journal of Health Economics*, 42:174–185, 2015.
- UNICEF-INJU-MIDES. Informe del panel de juventudes ENAJ 2018–2022. Technical report, MIDES, INJU and UNFPA, Montevideo, 2023.
- Courtney Harold Van Houtven, Norma B. Coe, and Meghan M. Skira. The effect of informal care on work and wages. *Journal of Health Economics*, 32(1):240–252, 2013.
- World Health Organization. From early detection to care: How Uruguay is transforming suicide prevention through real-time surveillance. WHO Feature Story, 2024.