



The marginal cost of mortality risk reduction: Evidence from housing markets[☆]

Kelly C. Bishop^a, Nicolai V. Kuminoff^{a,b}, Sophie M. Mathes^c, Alvin D. Murphy^{a,*}

^a Arizona State University, United States of America

^b NBER, United States of America

^c FGV EPGE, Brazil

ARTICLE INFO

Keywords:

Housing markets
Sorting
Hedonic
Mortality risk

ABSTRACT

We provide the first evidence on the rate at which spatial variation in all-cause mortality risk is capitalized into US housing prices. Using a hedonic framework, we recover the annual implicit cost of a 0.1 percentage-point reduction in mortality risk among older Americans and find that this cost is less than \$3453 for a 67 year old and decreasing with age to less than \$629 for an 87 year old. These estimates, while similar to estimates from the market for health care, are far below comparable estimates from markets for labor and automobiles, suggesting that the housing market provides an alternative, substantially cheaper channel for reducing mortality risk. We find this conclusion to be robust to a wide range of econometric model specifications, including accounting for associated expenditures on property taxes and the physical and financial costs of moving.

1. Introduction

People can modify their life expectancies by choosing where to live. This stylized fact is supported by causal evidence from economic studies and widely reported by the popular press.¹ While the precise mechanisms are not fully understood, the perception that mortality risk varies causally across residential locations, combined with the fact that people are free to move across locations, implies that mortality risk may be capitalized into housing prices. However, the capitalization rate may be dampened by information frictions and migration costs. This raises an important question: How much, if anything, do Americans need to pay via the housing market to reduce their mortality risk by moving?

This paper uses hedonic property-value methods to develop the first national evidence on the rate at which the US housing market capitalizes location-based mortality risk. We focus on the age group with the highest mortality: senior citizens. We start by following the [Finkelstein et al. \(2021\)](#) approach to measuring seniors' location-specific causal mortality risk. Specifically, we apply the identifying assumptions and

econometric methods from that paper to Medicare data (describing 7.2 million Americans aged 65 and above) to recover estimates of age- and location-specific all-cause mortality risk. We merge these causal mortality-risk estimates with national data describing house values, house characteristics, and location-specific amenities.

Our hedonic estimation strategy is designed to address two threats to identifying the capitalization of causal mortality risk. The first threat is that our estimates for causal mortality risk are likely to embed measurement error. To address this, we employ an Instrumental-Variables (IV) approach and instrument for causal mortality risk using raw, population-based measures of location-specific empirical mortality reported by the U.S. Centers for Disease Control and Prevention. As these empirical measures are simple population statistics, they are unlikely to embed measurement error. The second threat to identification is that mortality is likely to be correlated with unobserved determinants of housing prices. This threat applies to both the causal and empirical

[☆] We are grateful for insights and suggestions from Joe Aldy, Glenn Blomquist, Greg Caetano, Tatyana Deryugina, Catie Hausman, Matt Kahn, Jonathan Ketcham, Chris Parmeter, Jaren Pope, Edson Severnini, Aradhya Sood, and seminar audiences at Insper University, University of Alabama, University of California, Irvine, University of California, Santa Barbara, University of Queensland, University of São Paulo, FGV EESP, University of Southern California, University of Wisconsin - Madison, the ASSA Annual Meetings, the UEA North American Meeting, the AERE Annual Meeting, the AREUEA Annual Meeting, anonymous referees, and the editor, Stephan Heblich. This study analyzed secondary data on human subjects under Arizona State University IRB ID STUDY00001990. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

* Corresponding author.

E-mail address: alvin.murphy@asu.edu (A.D. Murphy).

¹ See, for example, [Weintraub \(2014\)](#), [Ferrari \(2017\)](#), [Ducharme and Wolfson \(2019\)](#) and [Ansari \(2022\)](#).

measures of mortality. In particular, some sources of mortality risk may be separately capitalized into housing prices through their effects on the quality of life.² We mitigate this threat by controlling for an extensive list of amenities suggested by the literature on measuring spatial variation in the quality of life (Roback, 1982; Albouy et al., 2016; Diamond, 2016). These controls include measures of climate, pollution, crime, transportation and other local public goods, in addition to state fixed effects. Thus, our estimates for the housing costs of reducing mortality risk are purged of differences in housing prices that may be explained by observed housing characteristics and salient amenities.

The residual variation in all-cause mortality risk that we use for identification retains location-specific sources of risk that are less likely to be capitalized into property values through channels other than mortality risk. The hedonic property value literature is replete with examples of risks to life and health that would be unlikely to directly affect quality of life (Banzhaf et al., 2019). Examples include unexplained cancer clusters (Davis, 2004), water contamination (Muehlenbachs et al., 2015; Christensen et al., 2023), toxicity of smoke emitted by facilities that handle hazardous chemicals (Banzhaf and Walsh, 2008), lead paint (Billings and Schnepel, 2017), and soil contamination (Ma, 2019). Capitalization of mortality risk may also be driven by the steady stream of news stories about spatial variation in longevity (Weintraub, 2014; Ferrari, 2017; Ducharme and Wolfson, 2019; Ansari, 2022).

Our approach to controlling for amenities raises two potential concerns. First, the amenities could be “bad controls” in the sense that they absorb too much relevant spatial variation in mortality risk. Second, the amenity controls could be insufficient to overcome confounding by latent amenities. We resolve both concerns by adapting and extending partial identification methods from Altonji et al. (2005), Banzhaf and Smith (2007), Nevo and Rosen (2012), Altonji et al. (2015), and Oster (2019) to derive upper bounds on the housing market capitalization of mortality risk using an assumption that capitalization of observed amenities is informative about capitalization of unobserved amenities. We test this assumption indirectly by examining how our IV estimates evolve as we incrementally add amenities in random order, generating over a quarter million different specifications. This analysis suggests that our amenity controls preserve ample identifying variation in mortality risk and that our IV estimates are consistent upper bounds on the actual housing market capitalization of mortality risk.³ This upper bound interpretation underscores the policy-relevance of our findings because our estimates are far below the most comparable estimates from the markets for labor and automobiles (Smith et al., 2004a; Aldy and Viscusi, 2008; Rohlfes et al., 2015; O’Brien, 2018; Banzhaf, 2022).

Our main IV results yield an estimated upper bound of \$1474 (year 2010 dollars) in annual housing expenditure at the household level to reduce annual mortality risk at age 77 by 0.1 percentage points (pp).⁴ A 0.1 pp reduction in annual mortality risk is approximately equal to the within-state standard deviation of mortality risk across locations. We find this implicit cost to be declining in age, from a high of \$3453 at age 67 to a low of \$629 at age 87. This is due to the fact that the spatial variation in mortality risk is increasing in age while the spatial variation in housing prices is not varying in age, making the implicit cost of mortality-risk reduction lowest among the oldest individuals.

Our estimates for the cost of reducing mortality risk provide information that can be shared with seniors or used to inform public policy. Numerous federal policies have incentivized migration (Jia et al., 2023), e.g., Freddie Mac and Fannie Mae’s support of the mortgage

² For example, air pollution increases mortality risk and also creates haze that impairs visibility (Banzhaf and Walsh, 2008; Deryugina et al., 2019; Kahn, 2004; Smith et al., 2004b).

³ This analysis is also suggestive that the housing market capitalization of mortality risk is likely close to our estimated upper bound.

⁴ This would translate to a per-person cost if households were comprised of single individuals. If multiple individuals live in a single housing unit this would reduce the per-person cost as discussed in Section 5.

market, the Moving-to-Opportunity experiment, IRS tax deductions for moving expenses, transferable unemployment insurance. Our findings suggest that determining the marginal value of public funds requires understanding how these policies affect migration, how migration affects mortality, and how much movers have to pay to reduce their mortality risks conditional on other amenities.

Our study adds to the literature estimating the cost of reducing mortality risk. This literature has mainly focused on estimating the wage compensation for undertaking a higher risk of on-the-job death (Viscusi and Aldy, 2003; Costa and Kahn, 2004; Cropper et al., 2011; Kniesner et al., 2012; Lee and Taylor, 2019; Evans and Taylor, 2020), with prevailing estimates exceeding \$6 million per statistical death among workers aged 60 to 65 (Smith et al., 2004a; Aldy and Viscusi, 2008; Banzhaf, 2022). To compare our findings with those figures, we rescale our main estimates to measure the aggregate housing expenditure required to avoid one statistical death at age 67, and find the implied figure to be at or below \$1.3 million. Our upper bound estimates are also an order of magnitude lower than those from studies that have estimated the cost of reducing mortality risk among older drivers via automobile safety features (Rohlfes et al., 2015; O’Brien, 2018). Our findings are much closer to estimates for the cost of reducing premature mortality among the general population of Americans over age 65 via medical spending (Hall and Jones, 2007; Doyle et al., 2015; Huh and Reif, 2017; Ketcham et al., 2022). For example, Hall and Jones (2007) find that the medical cost of avoiding an additional fatality is approximately \$1 million among people over age 65. Thus, our findings suggest that the marginal cost of reducing mortality risk in the housing market is much smaller than in the markets for labor and automobiles, and closer to the market for health care.

Our study also adds to the hedonic property-value literature on amenity capitalization. Portney (1981) first estimated the capitalization of mortality risk in a study of air pollution in Pennsylvania. Subsequent studies estimated capitalization of other specific mortality risks in specific areas, such as lead exposure in North Carolina (Billings and Schnepel, 2017), violent conflict in Northern Ireland (Besley and Mueller, 2012), and a cancer cluster in Nevada (Davis, 2004). Our study adds to this literature by providing the first nationwide analysis of mortality risk, the first analysis of mortality risk among senior citizens, and the first analysis to focus on a broadly inclusive measure of causal mortality.⁵

Finally, it is important to note that we make no claims about the demand for mortality risk reduction. We refrain from a revealed-preference interpretation of capitalization effects because it is unclear what, exactly, households believe about their ability to adjust their mortality risk by moving. Similarly, we take no stance on the relative importance of determinants of location decisions that would condition a welfare interpretation of our results, such as tastes for longevity, moving costs, location ties based on earlier-in-life decisions, or where children and grandchildren live. We leave the important question of how capitalization effects relate to private and social benefits of reducing mortality risk to future research.

2. Background: Capitalization of mortality risk

Life expectancy varies substantially across the US.⁶ For example, Dwyer-Lindgren et al. (2017) find that life expectancy at birth

⁵ By pricing residential mortality risk, our study also relates to literature on how location affects health, wealth, and human capital (Aliprantis and Richter, 2020; Barreca et al., 2015; Bayer et al., 2008; Caetano and Macartney, 2021; Card et al., 2022; Chetty and Hendren, 2018a,b; Couillard et al., 2021; Cutler and Glaeser, 1997; Deryugina and Molitor, 2020, 2021; Deschenes, 2014; Graff Zivin and Neidell, 2013; Kahn, 2004; Kling et al., 2007).

⁶ Researchers measure local life expectancy by aggregating location data from death certificates. A common source of these data is the U.S. Centers for Disease Control and Prevention (CDC), which reports annual county-specific mortality rates based on the universe of death certificates for U.S. residents. We describe these data in the next section.

differs by as much as twenty years between US counties. This range is striking and its precise causes are unknown. Unsurprisingly, life expectancy is strongly associated with income, education, and health behaviors (Case and Deaton, 2015; Chetty et al., 2016; Currie and Schwandt, 2016). These associations suggest that spatial variation in life expectancy might be explained by residential sorting based on socioeconomic status. At the same time, socioeconomic status during adulthood is believed to be affected by residential location during childhood (Chetty and Hendren, 2018a,b). These dynamics make it challenging to measure precisely how residential location modifies mortality risk. Recent studies have addressed this challenge by focusing on migrants. For example, Deryugina and Molitor (2021) and Finkelstein et al. (2021) analyze migration and mortality among the over-65 Medicare population and provide causal evidence that migrants' destination locations modify their remaining life expectancies.

While the literature suggests that older Americans can modify their life expectancies by moving, the exact production function for location-specific causal mortality risk remains unknown. There are numerous environmental and urban amenities that may contribute. Some are well known. For example, public opinion polls indicate that people tend to worry about local air pollution and water contamination (Lin et al., 2021). Since people can avoid these and other salient health risks by moving, one might expect location-specific risks to be capitalized into property values. This is supported by evidence on the capitalization of quasi-random sources of localized variation in air pollution (Banzhaf and Walsh, 2008) and water contamination (Muehlenbachs et al., 2015; Christensen et al., 2023), as well as other salient risks such as crime (Besley and Mueller, 2012), cancer (Davis, 2004), lead paint (Billings and Schnepel, 2017), and soil contamination (Ma, 2019).

However, the cumulative capitalization of local health risks may not reveal what the average person is willing to pay to reduce their own mortality risk due to market frictions such as moving costs and incomplete information. Moving costs can attenuate the capitalization rate of an amenity by preventing inframarginal households from adjusting to changes in amenity levels or new information (Bayer et al., 2009). Similarly, if households have incomplete information about an amenity then, in general, their location choices cannot be expected to reveal their willingness to pay (Leggett, 2002). In particular, Pope (2008b) and Kumbhakar and Parmeter (2010) show how incomplete information can create a wedge between buyers' marginal willingness to pay for an amenity and the equilibrium hedonic price function gradient.⁷ This conceptual framework is supported by evidence that capitalization rates for local amenities are attenuated by inattention to publicly available information about those amenities (Pope, 2008a,b,c; Ma, 2019).⁸ Providing information about either local amenities or the capitalization rate of local amenities could potentially reduce this wedge as households respond and would likely increase the marginal cost of the amenity.

In summary, the prior literature has shown that some health risks are capitalized into property values, and it has also shown that capitalization rates can understate households' willingness to pay due to moving costs and incomplete information. These frictions may be quantitatively important in our setting for at least three reasons. First, the scientific evidence on our variable of interest – the overall location-specific causal mortality risk for seniors – was published after the households in our data made location decisions. It seems unlikely that households would have fully anticipated the findings in Finkelstein

et al. (2021). Second, while housing prices may reflect spatial variation in some of the environmental health risks that contribute to the findings in Finkelstein et al. (2021), there is substantial evidence that households may not be fully attentive to public information on all risks (Pope, 2008a,b,c). Finally, households' understanding of the spatial variation in causal mortality risk may be complicated by the media's extensive coverage of associative studies of life expectancy with headlines such as “Lifespan More to Do with Geography than Genetics” (Weintraub, 2014) and “Your Zip Code Might Determine How Long You Live – and the Difference Could be Decades” (Ducharme and Wolfson, 2019).⁹ Against this background, we focus on estimating the capitalization effects of causal mortality risk and refrain from interpreting these effects as measures of household willingness to pay.

3. Data

3.1. Locations

We begin by defining location at the level of commuting zones (CZs), as defined by the U.S. Economic Research Service using 2000 Census data, for the 48 coterminous US states. Each CZ is a cluster of counties that approximate a local labor market and is similar in size to a metropolitan area. There are 709 distinct CZs in our data. To deal with the issue of sparsely-populated CZs, we follow the procedure in Finkelstein et al. (2021) (FGW) which aggregates CZs into 536 aggregate CZs (which we continue to refer to as CZs for simplicity).

This CZ definition of location is well suited to measuring the housing-market capitalization of causal mortality risk.¹⁰ A much finer resolution of geography (e.g., a Census tract) would exacerbate any potential for measurement error, as individuals would likely spend a considerable amount of time outside of their home location, while a coarser resolution of geography (e.g., a state) would limit the scope to describe individuals' ability to adjust their mortality risk by moving reasonably short distances.¹¹

3.2. Location-specific empirical mortality rates

We use the “Multiple Cause of Death” data produced by the U.S. Centers for Disease Control and Prevention (CDC) as a measure of spatial variation in all-cause mortality. These data include annual county-level mortality rates for each integer age from 65 to 84. For ages 85 and over, the data are right-censored. The underlying data on deaths are derived from the population of death certificates, while county population sizes are derived from Census data.

We refer to the location-specific mortality rates as “empirical” as they simply describe the mortality rates that are observed within locations. Panel (a) of Fig. 1 shows the spatial variation in empirical mortality at age 77, which is approximately the average age in the over-65 population. The variation across locations is substantial. For example, moving from the 10th percentile to the 90th percentile in the distribution is associated with a 1.48pp increase in mortality risk over one year (a 42% percent increase).

To construct our empirical-mortality instrument, we synthesize these population-based data into a single IV, which we construct by taking a weighed average over a CZ's empirical mortality rates at each year, county, and integer age, weighting by Census-based measures


⁷ The model of incomplete capitalization in these studies is an example of the broader literature on how market frictions condition the interpretation of equilibrium hedonic price functions and structural parameters describing market primitives (Harding et al., 2003).

⁸ See Hausman and Stolper (2022) for a review of the literature on information frictions and amenity capitalization.


⁹ Ferrari (2017) and Ansari (2022) provide additional examples.

¹⁰ National hedonic studies typically partition the country into similar geographies such as metropolitan areas (Bayer et al., 2009), counties (Blomquist et al., 1988), or public-use microdata areas (Albouy et al., 2016).


¹¹ In cases where there are multiple counties within a commuting zone, using a county may lead to measurement error if individuals spend time outside of their home county. An additional benefit of using CZ is that it allows us to compare our mortality estimates with those found in FGW.



(a) Empirical Annual Mortality Rates at Age 77 by Commuting Zone



(b) Causal Annual Mortality Rates at Age 77 by Commuting Zone



(c) Average Annual Log Rents by Commuting Zone

Fig. 1. Spatial variation in empirical mortality, causal mortality, and housing rents.

for U.S. population shares by age.¹² We prefer this construction of the instrument because it combines all of the available CDC data, but we also show in Appendix D that our results are essentially unchanged if we replace our preferred instrument with CDC mortality rates observed at a single integer age (e.g., 65, 75, or 85+).

Importantly, the variation in empirical mortality rates cannot be interpreted as causal. While some of the variation may be caused by

placed-based amenities that affect mortality risk, some of the variation may arise from healthier and wealthier people sorting themselves into higher amenity areas and living longer for reasons independent of their residential locations.

3.3. Location-specific causal mortality rates

To disentangle location-based inputs to mortality risk from residential sorting, we follow FGW both in applying the estimator developed in that paper and in describing the resulting measures as providing the “causal” effect of location on mortality. The analysis in that paper

¹² Applying the same national weights to each CZ avoids compositional bias that would be introduced if we were to instead weight by CZ-specific population shares.

uses Medicare data describing people over the age of 65 to recover CZ-specific mortality risks while controlling for differences in age, race, gender, medical spending, and clinical diagnoses of chronic medical conditions. The identification strategy leverages the variation in survival rates among movers who originate from the same location but who move to different locations. A selection-correction procedure is used to account for any sorting based on unobserved health factors. We replicate the procedures described in that paper using administrative records for a 20% random sample of Medicare beneficiaries (7.2 million people) that we observe from 1999 through 2013 to estimate causal mortality for seniors aged 67 and older. Our results are very similar to the estimates reported in FGW and details are provided in Appendix A. As our measures of causal mortality are model-based estimates, they may differ from true causal mortality. We address this measurement error as part of the identification strategy discussed in Section 4.2.

Panel (b) of Fig. 1 shows the spatial variation in causal mortality at age 77. There is less spatial variation in causal mortality rates than in empirical mortality rates, speaking to the role of residential sorting; our estimates imply that moving from the 10th percentile to the 90th percentile in the distribution across locations causes a 0.28 pp increase in causal mortality risk over one year, compared with a 1.85 pp increase in empirical mortality risk. The standard deviation of mortality risk across locations is 0.13 pp.

We find substantial within-state variation in causal mortality rates across locations. For example, regressing age-77, location-specific causal mortality rates on a vector of state dummies yields an R^2 of 0.25 and a residual within-state standard deviation of mortality risk of 0.11 pp. These results highlight the scope for differences in location-specific mortality rates to be capitalized into property values, and, importantly, may inform policy based on the ability of older Americans' to lower their mortality risk through within-state moves.

3.4. Housing characteristics and prices

Data describing housing characteristics and prices come from the 2000 Census Integrated Public Use Microdata Series (IPUMS) 5% sample. In this sample, locations are defined as Public Use Microdata Areas (PUMAs), which are contiguous areas comprised of approximately 100,000 individuals each. The data contain a total of 2071 PUMAs which we aggregate into CZs.

For each of the 5.1 million properties in our sample, we observe the number of rooms, bedrooms, and units in the structure, whether the property has a kitchen and indoor plumbing, and the age of the structure. We additionally observe the self-reported value of the property, if owner-occupied, or the gross rent, if renter-occupied.¹³ Table B.1 reports summary statistics.

Panel (c) of Fig. 1 shows the spatial variation in the log of annual gross rents across locations. Fig. 1 suggests that mortality and log rents are correlated across space and the unweighted correlation coefficient between causal mortality at age 77 and log rents is -0.2 . This unconditional negative correlation may be driven by a causal relationship and/or by amenities that households value for reasons apart from mortality. Thus, we compile data describing a large set of location-specific amenities.

¹³ An alternative data source would be housing transactions datasets sold by commercial vendors. Transactions data have the advantage of typically including measures of square footage and the most recent sale price. However, transactions data are not nationally representative: they typically exclude renters as well as sale prices in states with non-disclosure laws. Due to these limitations, IPUMS data are typically preferred in national hedonic studies (Albouy et al., 2016; Diamond, 2016). Additionally, the IPUMS data report when each house was last sold enabling us to condition on recent sales whose reported values are likely to be similar to their transaction prices. See Section 5.4 for details.

3.5. Additional location-specific amenities

We compile data on location-specific amenities from several sources. We begin by collecting data on the county-level amenities used in Diamond (2016). In that paper, the variables describe urban counties. For our analysis, we measure amenities for rural counties, too, and include additional climate variables for all counties. We aggregate all county-level amenities to the level of the CZ, weighting by population.

Our final set of 18 amenities is comprised of measures of summer temperature, winter temperature, precipitation, fine particulate air pollution (PM_{2.5}), ozone concentrations, violent crime, property crime, student-teacher ratios in local public schools, interstate highway mileage, urban arterial mileage, number of urban rail stops, unemployment rate, share of residents with college degrees, and per-capita measures of government spending on parks, government spending on schools, number of movie theaters, number of restaurants and bars, and number of apparel stores.¹⁴ Appendix C documents the data compilation and Table C.1 reports summary statistics.

4. Econometric model

4.1. A hedonic price function of mortality risk

To estimate the effect of causal mortality on house prices, we specify Eq. (1) as a hedonic price function where the dependent variable, $\log p_{hj}$, describes the log price associated with occupying house h in location j . We follow Bayer et al. (2007) in pooling data on rents and property values by measuring price, p_{hj} , as annual gross rent, if renter-occupied, and property value, if owner-occupied. Controlling for owner-occupancy with the indicator own_h allows the coefficients on housing characteristics, x_h , location-specific amenities, x_j , and location-specific causal mortality risk measured at an arbitrary reference age, m_j , to be interpreted as annual measures of their implicit prices.¹⁵

$$\log p_{hj} = \alpha_1 + \alpha_2 own_h + \alpha_3 x_h + \alpha_4 x_j + \beta m_j + \varepsilon_{hj}. \quad (1)$$

Conditional on employing the Gompertz specification used in FGW, the reference age at which we measure mortality risk can be chosen without loss of generality. The Gompertz specification allows mortality risk to vary with location, age, and other characteristics:

$$m(\text{age}, j) = \exp(\delta \cdot \text{age} + \bar{\theta} + \gamma_j). \quad (2)$$

In this specification, δ is the estimated scaling parameter on age and $\bar{\theta}$ is the estimated index of other individual characteristics (race, gender, medical expenditures, and diagnoses of chronic medical conditions) that we scale to its national average. γ_j is a CZ fixed effect that captures how location contributes to mortality risk. It is the only component of $m(\text{age}, j)$ that varies across locations. m_j is then simply defined as $m(\text{age}^R, j)$, where the arbitrary reference age is denoted by age^R .

When embedded within the hedonic price function, the Gompertz specification used by FGW has two important implications for estimating the implicit cost of reducing mortality risk. First, causal mortality

¹⁴ This list excludes measures of local health care quality. Such measures are typically excluded from the set of amenities considered in empirical literature on quality-of-life and residential sorting (Roback, 1982; Albouy et al., 2016; Diamond, 2016). We follow the literature in excluding these measures from our main analyses. We do so because of their potential to be "bad controls" (Angrist and Pischke, 2009; Cinelli et al., 2022) that could absorb the effect we intend to measure. However, we also note that including such measures does not change our results substantially. Our main estimates decline by 10% if we extend the set of amenity covariates to include the rate at which Medicare patients are discharged from local hospitals after being admitted for ambulatory care sensitive conditions, a common measure of hospital quality.

¹⁵ Our main econometric specification allows α_1 and α_2 to differ by US state.

for any given age is simply a scaled version of causal mortality for any other age.¹⁶ This means that the price regression estimated for a reference age, shown in Eq. (1), can be used to calculate the corresponding coefficient on mortality risk for any other age as:

$$\beta_{age} = \beta \cdot \exp(\delta \cdot (age^R - age)). \quad (3)$$

Second, the marginal cost of reducing mortality risk will decrease in age if mortality risk increases in age (i.e., $\delta > 0$) and housing price decreases in mortality risk (i.e., $\beta < 0$).¹⁷ This model feature is intuitive. While mortality risk increases in age, all individuals face the same menu of housing prices, regardless of age.

We define the marginal cost of mortality risk reduction (MCMRR) at a given age as:

$$MCMRR_{age} = -\beta_{age} \cdot \bar{p}. \quad (4)$$

The MCMRR is simply the derivative of the hedonic price function in Eq. (1) with respect to mortality risk, evaluated at an age-specific mortality risk and the mean annual housing cost over all houses, $\bar{p} = 1/H \sum_{h=1}^H p_{hj} / \exp(\alpha_2 \cdot own_h)$.¹⁸ It measures the annual housing cost to a household of a marginal reduction in mortality risk at a given age. The main challenge in identifying the MCMRR is to develop a consistent estimator for β .¹⁹

4.2. Identifying an upper bound on the MCMRR

Our estimation approach is designed to address two threats to identifying the hedonic parameter on mortality risk, β . The first threat is that our estimates of location-specific mortality effects are likely to be measured with error. Formally, this may be written as: $\hat{m}_j = m_j + \xi_j$, where \hat{m}_j denotes our measure for location-specific causal mortality risk derived using the FGW estimator and treated as data in the price regression. The second threat is that our estimates of location-specific mortality effects may be driven, in part, by unobserved amenities that are directly capitalized into housing prices because they simultaneously affect both the quality and the quantity of life. Formally, the set of location-specific amenities can be partitioned as: $x_j = [x_j^o, x_j^u]$, where o and u denote the subsets of observed and unobserved amenities. Analogously, $\alpha_4 = [\alpha_4^o, \alpha_4^u]$. Eq. (5) rewrites Eq. (1) to highlight these two potential sources of endogeneity.

$$\log p_{hj} = \alpha_1 + \alpha_2 own_h + \alpha_3 x_h + \alpha_4^o x_j^o + \beta \hat{m}_j + v_{hj}, \quad (5)$$

where $v_{hj} = \varepsilon_{hj} + \alpha_4^u x_j^u - \beta \xi_j$.

We interpret ε_{hj} as idiosyncratic noise and focus on potential threats to identification from unobserved variables, x_j^u , and measurement error, ξ_j .²⁰

¹⁶ $m(age, j) = m(age', j) \cdot \exp(\delta \cdot (age - age')) \forall age, age'$.

¹⁷ Like FGW, we find that δ is positive ($\hat{\delta} = 0.0977$) and our estimates for β are universally negative across numerous specifications.

¹⁸ While hedonic studies typically evaluate price functions at average house prices, Eq. (4) can be evaluated at any house price. For example, if we were to evaluate it at the price that the average individual pays for housing, conditional on age, that measure would embed the effects of sorting along the price function; e.g., older households choosing to locate in more expensive locations where life expectancy is higher.

¹⁹ [Albouy \(2009\)](#) shows that geographically varying tax rates may also affect households' location decisions. Although labor market participation is low among seniors, their fixed incomes may be subject to geographically varying income taxes. However, the within-state variation in taxes on these fixed incomes is minimal. Within-state variation in property taxes could have a more substantial effect on the real cost of housing paid by seniors ([Shan, 2010](#)). In Section 5.4 we show that accounting for within-state variation in property taxes increases our main estimates for the cost of mortality risk reduction by approximately 5%.

²⁰ ε_{hj} captures the standard modeling errors that are often thought to be relatively innocuous in hedonic price function estimation, such as measurement error in prices, functional form mis-specification, and omitted architectural details.

We address the omitted-variable and measurement-error threats separately. First, we use a population-level measure of empirical mortality, z_j , to instrument for causal mortality. As the instrument, empirical mortality, is simply a population statistic, we have no reason to expect it to embed measurement error. Our endogenous variable of interest, causal mortality, is effectively a selection-corrected version of empirical mortality and, in contrast to empirical mortality, we allow for causal mortality to be measured with error. In particular, measurement error, ξ_j , is likely introduced through the inevitable modeling assumptions used in the selection-correction procedure of the causal-mortality estimator. We assume that this measurement error in estimated causal mortality is not systematically related to empirical mortality rates²¹:

$$cov(z_j, \xi_j) = 0. \quad (6)$$

This IV strategy addresses measurement error, but remains vulnerable to confounding from $\alpha_4^u x_j^u$. The concern is that empirical mortality may be correlated with the composite error if important amenities that determine price cannot be observed and are correlated with empirical mortality, so that $cov(z_j, \alpha_4^u x_j^u) \neq 0$. Specifically, if unobserved amenities are correlated with the instrument, we would expect the net effect of this correlation to be weakly negative, i.e.,

$$cov(z_j, \alpha_4^u x_j^u) \leq 0. \quad (7)$$

Intuitively, the bundle of unobserved amenities that increases the quality of life (and therefore housing prices) is, on net, likely to be associated with increased longevity. Even if unobserved amenities have no direct effects on longevity, the association in Eq. (7) is likely to arise from healthier and wealthier people sorting into higher-amenity locations and living longer.

It is worth noting that Eq. (7) permits the bundle of unobserved amenities to include a subset of amenities that increase the quality of life (and housing prices) and decrease longevity, e.g., state laws that increase access to controlled substances and gambling. This is because Eq. (7) only restricts the sign of the *net* effect of the bundle of unobserved amenities on longevity (i.e., $\alpha_4^u x_j^u$), not the sign of the effect of any particular amenity in that bundle.

Eq. (7) is not directly testable. However, it is likely to hold as it matches two stylized facts from the residential sorting literature: (i) wealthier households sort into higher amenity areas and (ii) amenities that increase housing prices tend to be positively correlated over space because their production functions embed location-specific features such as climate, geography, and the property tax base ([Banzhaf and Walsh, 2008](#); [Kuminoff et al., 2013](#)). Further, following [Altonji et al. \(2005\)](#) and [Oster \(2019\)](#), Eq. (7) can be indirectly tested under the assumption that capitalization of observed amenities is informative about capitalization of unobserved amenities in the sense that the impact on $\hat{\beta}$ from omitting unobserved amenities has the same sign as the impact of omitting observed amenities:

$$\frac{\hat{\beta}(W, x_j) - \hat{\beta}(W, x_j^o)}{\hat{\beta}(W, x_j^o) - \hat{\beta}(W)} \geq 0, \quad (8)$$

where $W = [p_{hj}, own_h, x_h, z_j]$. The numerator is the unmeasurable effect of adding x_j^u as covariates in the hedonic regression. It equals zero in the special case where all relevant amenities are observed: $x_j^o = x_j$. The denominator is the measurable effect on $\hat{\beta}$ from adding x_j^o .

Taken together, Assumptions (6) and (8) provide us with the direction of any inconsistency caused by omitted amenities and allow us to apply the intuition of the ‘‘imperfect IV’’ strategy from [Nevo and Rosen \(2012\)](#). The key insight from [Nevo and Rosen \(2012\)](#) is that even when an instrument is correlated with the error, a bound is identified when

²¹ A violation of this assumption would require that the measurement error in the estimates of m_j is systematic, which would violate the identifying assumptions in FGW.

the correlation between the instrument and the error can be signed. In our case, the instrument plausibly identifies an upper (lower) bound on the MCMRR (β).²²

In summary, our strategy for addressing potential confounding by unobserved location-specific amenities is to control for an extensive set of observable amenities and to interpret our MCMRR estimates based, in part, on how they evolve as we expand the set of amenity controls.

5. Results

5.1. Estimates of the marginal cost of mortality risk reduction

Table 1 reports our estimates for the MCMRR from five different specifications of the price function. For each specification, we report estimates for ages 67, 72, 77, 82 and 87 followed by robust standard errors clustered at the CZ level.²³ All five specifications control for physical house characteristics, but differ in the steps taken to mitigate confounding.

The estimates in Column (1) are from an OLS regression that excludes amenities. We report the MCMRR for a 0.1 pp reduction in the annual probability of death among a given age group. This reduction is roughly equivalent to the within-state standard deviation of causal mortality risk across locations for people at age 77. The results show that the marginal reduction in mortality risk among 77-year-olds is associated with a \$633 increase in annual housing costs. This associative measure declines with age from \$1482 among 67-year-olds to \$270 among 87-year-olds.²⁴ We use the \$633 value for the 77-year old group as a benchmark when comparing with other specifications, as 77 is close to the mean age among the 65-and-over population. The results in Column (1) do not have a causal interpretation as they embed potential biases from measurement errors in mortality risk and omitted amenities, and the net direction of these biases is ambiguous. We address these issues incrementally.

To address measurement error in causal mortality risk, we use our population-based empirical-mortality variable as an instrument in Column (2). The results show that, all else constant, moving from OLS to IV increases the estimates by a factor of 6, yielding a MCMRR of \$4015 at age 77. This six-fold increase is large in an absolute sense, but it is small relative to prior estimates for the effects of using IVs to address attenuation bias from measurement error in other amenities such as air pollution (Schlenker and Walker, 2016; Deschenes et al., 2017; Deryugina et al., 2019; Bishop et al., 2023b). In addition, the six-fold increase captures the combined effect of the instrument reducing the likelihood of the estimate being below the true effect due to measurement error in causal mortality and increasing the likelihood of the estimate being above the true effect due to healthier and wealthier people sorting into higher-amenity areas and living longer. The next two columns take steps to disentangle these mechanisms.

In Column (3), we narrow our focus to within-state variation in causal mortality by adding state dummies along with interactions of the state dummies and the owner-occupancy indicator. This reduces the implied MCMRR at age 77 to \$3216. Adding state dummies sharpens the identification strategy in three ways. First, the state dummies absorb the price effects of between-state variation in omitted amenities that are correlated with mortality risk. Second, the interactions absorb any

²² Analogously, in a model with no amenity controls, the instrument can identify a less sharp bound on the MCMRR.

²³ We construct standard errors based on sampling variation in the price regression only, as we treat the first stage estimates of \hat{m} as incorporating measurement error. We report the underlying estimates for β and \bar{p} in Appendix Table D.1, and report MCMRR estimates for intermediate ages in Appendix Table D.3.

²⁴ As discussed in Chetty et al. (2016), FGW, and Section 4.1, the rate of decline is determined by the functional form of the Gompertz specification in addition to the estimate of how mortality increases in age.

between-state variation in the user-cost of housing that is correlated with mortality risk (Poterba, 1984). Finally, focusing on within-state variation in prices and mortality risk reduces the potential concern that moving costs and information frictions may limit the extent to which housing markets capitalize spatial variation in mortality risks.

The estimator in Column (3) remains vulnerable to confounding from within-state sorting on amenities. For example, wealthier people with longer life expectancies may tend to locate in higher-amenity locations within their home states. This sorting behavior would impart an upward bias to our MCMRR estimator as we would expect omitted amenities to be positively correlated with housing prices while negatively correlated with empirical mortality.

We address this threat by augmenting the IV estimator in Column (3) to add the 18 amenity covariates. Adding these covariates reduces the age-77 MCMRR to \$1474.²⁵ This decline is consistent with the intuition that people live longer in locations where housing is more expensive due to amenity capitalization. We consider these results, shown in Column (4), to be our main MCMRR estimates.

5.2. Instrument validity

Under Assumptions (6) and (8), the results in Column (4) of Table 1 provide consistent estimates of an upper bound on the MCMRR.²⁶ While an exclusion restriction, such as $cov(z_j, a_{k,j}^u) = 0$, would be sufficient to point identify the MCMRR, we require only the weaker condition of a directional restriction on the sign of the net effect of the bundle of unobserved amenities on longevity, $cov(z_j, \alpha_{k,j}^u x_j^u) \leq 0$, to identify an upper bound.²⁷ To provide support for our identifying, directional-restriction assumption, we examine how our MCMRR estimates evolve as we add amenity covariates incrementally.

We show the results of this analysis in Fig. 2. The solid curve shows how our MCMRR estimates at age 77 change as we incrementally add amenity covariates to move from the specification shown in Table 1 Column (3) that excludes all amenity covariates (denoted by the square) to the specification shown in Column (4) that includes all of our 18 amenity covariates (denoted by the triangle). To avoid sensitivity to the order in which amenities are added, we estimate models for all 262,144 possible combinations of the 18 amenity covariates. Each point on the solid curve shows the mean MCMRR (measured on the left vertical axis) estimated over all models that use the number of amenity covariates shown on the horizontal axis.

This exercise examines what happens when we begin by assuming all amenities are unobserved and then treat various combinations of the 18 amenities as observed, and shows how our estimates evolve, on average, as amenity covariates are added. The mean MCMRR declines monotonically in the number of amenities, as expected. We view this as supporting evidence for the key identifying assumption, Eq. (8). Under this assumption, the insights of Nevo and Rosen (2012) apply, and we consistently estimate an upper bound for the MCMRR.

Importantly, the curvature, as well as the slope, of the solid curve in Fig. 2 is potentially informative about the MCMRR. In particular, the mean MCMRR declines at a decreasing rate; the slope is close to zero by

²⁵ This estimate is 7.7 times larger than the estimate that we obtain from an OLS analog to this specification, reinforcing the importance of addressing attenuation bias from classical measurement error.

²⁶ The results in Column (4) of Table 1 would provide consistent point estimates under the assumption that our amenity covariates span the set of amenities that both determine housing prices and are correlated with empirical mortality. Although we do not make this assumption, we provide evidence below that our amenity covariates are sufficiently rich to suggest that the MCMRR is likely to be close in magnitude to our estimated upper bound.

²⁷ A considerably stronger assumption that would also be sufficient to identify an upper bound is that each unobserved amenity, k , is negatively correlated with the instrument, i.e., $\max_k \{cov(z_j, \alpha_{k,j}^u x_{k,j}^u)\} \leq 0$. However, this stronger assumption is unnecessary.

Table 1
The marginal cost of mortality risk reduction.

	(1)	(2)	(3)	(4)	(5)
Marginal cost of mortality risk reduction, age 67	1482 (555)	9405 (2893)	7533 (1592)	3453 (729)	3023 (677)
Marginal cost of mortality risk reduction, age 72	967 (362)	6145 (1891)	4922 (1041)	2257 (477)	1974 (442)
Marginal cost of mortality risk reduction, age 77	633 (237)	4015 (1235)	3216 (680)	1474 (312)	1290 (289)
Marginal cost of mortality risk reduction, age 82	413 (154)	2624 (807)	2102 (444)	964 (204)	843 (189)
Marginal cost of mortality risk reduction, age 87	270 (101)	1715 (527)	1373 (290)	629 (133)	551 (123)
1st stage coefficient on instrument		0.079 (0.019)	0.084 (0.013)	0.073 (0.011)	0.074 (0.012)
1st stage F-statistic		17.1	42.4	44.3	40.9
State dummies			x	x	x
Amenity covariates				x	x
Recent moves only (last 5 years)					x
Clustering (number of CZs)	536	536	536	536	536
Number of houses	5,118,639	5,118,639	5,118,639	5,118,639	2,413,212

Note: The table reports estimates for the annual housing price of reducing the annual risk of death by one tenth of a percentage point for each age group. Estimates are reported in 2010 dollars. All specifications include housing covariates. Robust standard errors are clustered by CZ.

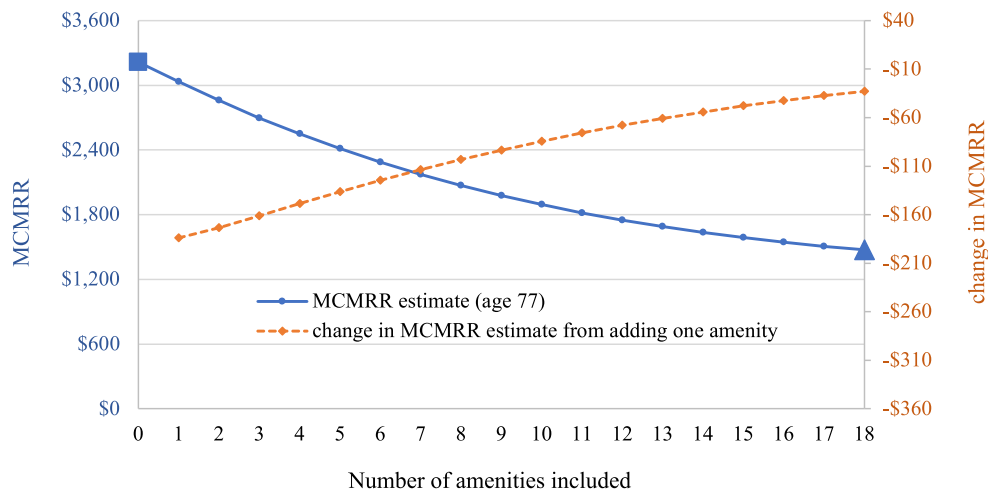


Fig. 2. Sensitivity of results to amenity covariates across 262,144 IV regressions.

the time we add the final amenity. To further illustrate this point, the dashed curve plots the change in mean MCMRR as we incrementally add amenities, i.e., the gradient of the solid curve. The size of the change is measured on the right horizontal axis. As we randomly add the first amenity to the model, the mean MCMRR declines by \$179. However, adding the last observed amenity reduces the mean MCMRR by only \$34. The curvature of the mean MCMRR function suggests that if one were able to estimate the MCMRR using all potential amenities, then that estimate would likely be relatively close to our estimated upper bound.

We use this exercise to incorporate and build upon the insights of Altonji et al. (2005), Banzhaf and Smith (2007), and Oster (2019). Broadly speaking, the approach taken in Oster (2019) would be to project out the mean MCMRR (the solid curve in Fig. 2) via linear extrapolation until the R^2 reaches its conceptual maximum. Under the assumption that observables are at least as important as unobservables, this would yield an estimated lower bound on the MCMRR. In our case, however, the R^2 is not informative as we use IV. Alternatively, by leveraging the curvature in addition to the gradient of the MCMRR function, we capture the key insight of Oster (2019) without relying on an R^2 . The gradient and curvature suggest that the MCMRR is likely close to our estimated upper bound.

Given a functional form assumption, one could project the curves in Fig. 2 to predict what would happen if one were to add more hypothetical unobserved amenities as controls. A log-quadratic function fits the mean MCMRR curve almost perfectly and reaches a minimum at 27 amenities. This corresponds to an MCMRR of \$1323, which is 90% of our upper bound estimate.²⁸ Under the assumption that Fig. 2 is informative about what happens to the estimated MCMRR as controls are added in random order and the assumption that observed amenities are at least as important as unobserved amenities, \$1323 can be interpreted as an estimated lower bound on the MCMRR. When interpreting our results, we choose to focus on the estimated upper bound for two reasons. First, it requires substantially weaker assumptions than the lower bound. Second, the upper bound is arguably the more important economic result, as it provides evidence that the cost of reducing mortality through the housing market is relatively low.

²⁸ As extrapolation far beyond the range of the data is reliant on functional form assumptions, we view this exercise as suggestive that the MCMRR curve will become approximately flat with the addition of a relatively small number of additional amenity controls. While log-linear and polynomial functions fit the mean MCMRR curve poorly, a log-cubic function fits almost perfectly and yields a less-conservative lower-bound MCMRR of \$1406.

5.3. Interpretation of results

To compare our results with existing evidence on the implicit per-capita cost of reducing mortality risk, we first multiply our household-level MCMRR estimates by 1000 to measure the capitalization associated with a one-unit reduction in the probability of a death. We then divide by the average household size in our data, 2.59, in order to rescale MCMRR to the individual level.²⁹ In other words, we measure the annual housing expenditures needed to avoid one premature statistical death in expectation.³⁰

Applying this transformation to our main specification in Column (4) of Table 1 yields a cost of approximately \$1.3 million to avoid a statistical death among 67-year-old residents. This cost declines with age to approximately \$0.6 million at age 77 and \$0.2 million at age 87. We interpret these measures as upper bounds on the actual housing cost of reducing mortality risk, consistent with Assumptions (6) and (8) and the supporting evidence in Fig. 2.

Prior evidence on the cost of reducing mortality risk is largely based on the wage compensation for undertaking a higher risk of death on-the-job. Prevailing estimates from the hedonic wage literature range from \$6 to \$10 million per avoided fatality (Viscusi and Aldy, 2003; Costa and Kahn, 2004; Cropper et al., 2011; Kniesner et al., 2012; Lee and Taylor, 2019; Evans and Taylor, 2020). However, these estimates are derived from a labor force that is considerably younger than the population that we study. The closest that we can come to conditioning on age is to narrow our focus to people in their early-to-mid 60s. Our housing market estimate for 67-year-old residents is \$1.3 million, compared to labor market estimates for 60-to-65-year-old workers that exceed \$6 million (Smith et al., 2004a; Aldy and Viscusi, 2008; Banzhaf, 2022). For this age group, our results suggest that it is substantially cheaper to reduce mortality risk through the housing market compared to the labor market.

To provide a more direct comparison to the general population of senior citizens, we can compare our findings to a few studies that have estimated their costs of reducing mortality risk via automobile safety features and health care. In the automobile market, Rohlfs et al. (2015) uses a hedonic regression of vehicle prices and fatality rates to conclude that the median cost of avoiding a premature death, via airbags, is \$9 to \$11 million for the general population of drivers, and slightly higher for drivers over age 60.³¹ In the market for health care, Hall and Jones (2007) finds that the medical cost of avoiding an additional fatality is approximately \$1 million among people over age 65.³² Thus, our findings suggest that the marginal cost of reducing mortality risk in the housing market is much smaller than in the markets for labor and automobiles, similar to the market for health care in the late 60s, and smaller than the market for health care at older ages. These comparisons are reinforced by the upper bound interpretation of our estimator for capitalization effects.

²⁹ While we follow Davis (2004) in dividing by the average household size, our results can be easily rescaled for any household size.

³⁰ As noted above, we do not interpret this result as a welfare measure. Our lack of revealed preference interpretation differentiates our approach from wage-hedonic studies that typically interpret cost estimates as measures for the “value of statistical life”.

³¹ Similarly, O’Brien (2018) obtains an estimate of approximately \$9 million for drivers aged 65 to 85 based on a logit model of automobile choice that adjusts for age-specific driving intensity and age-by-automobile-specific fatality rates.

³² In Hall and Jones (2007), the estimated cost of reducing mortality risk is relatively insensitive to increasing age beyond age 60. For example, the cost among seniors aged 90 to 94 is just 7% lower than among those aged 60 to 64. Other studies that have focused on subsets of Medicare patients have reported average costs similar to those in Hall and Jones (2007) but, in some case, with steeper declines in age and/or morbidity (Doyle et al., 2015; Huh and Reif, 2017; Ketcham et al., 2022).

A caveat to these comparisons is that they ignore transaction costs, which may be higher in the housing market than in markets for labor, automobiles, or health care. To put housing transaction costs into perspective, Bieri et al. (2023) constructs measures for the physical and financial costs of moving between every pair of locations in the contiguous US. These costs include realtor fees, closing costs, costs for home-finding trips, and the costs of hiring a company to move personal belongings and cars.³³ Taking an average over the costs of moving between each pair of locations and converting to year 2010 dollars yields a one-time cost of approximately \$15,000. While these one-time costs are non-trivial, they do not change the conclusion that it is substantially cheaper to reduce mortality risk through the housing market compared to the labor and automobile markets.³⁴

5.4. Sensitivity analysis

While the IPUMS data that we use to measure property values are routinely used in national hedonic studies (Albouy et al., 2016; Diamond, 2016) a potential concern is that self-reported property values may be measured with error. We address this concern by repeating the estimation after restricting the sample to properties that changed occupant within the previous five years since recent migrants have made a costly decision to move and their experience provides them with the information needed to report their housing values with greater accuracy (Bajari and Kahn, 2005). Results of this sensitivity check are presented in Column (5) of Table 1 and show that our featured estimates decline by 12%.

Another potential concern is that our featured estimates could understate the full cost that homeowners would have to pay to reduce their mortality risk by moving because our measure of housing prices excludes property taxes, which could also be spatially correlated with mortality risk. This is particularly important to consider in light of evidence that property taxes influence older homeowners’ migration decisions (Shan, 2010). We address this concern by repeating the estimation after converting owner-occupied house values to annualized measures of the user-cost of housing. To make this conversion, we multiply the self-reported values of owner-occupied houses by location-specific user-cost rates from Bieri et al. (2023). These location-specific user-cost rates account for between-PUMA variation in local property taxes, as well as the ability to deduct property taxes and mortgage interest payments when filing federal income taxes. This modification to our dependent variable increases our featured estimates in Table 1 by just under 5%.³⁵

We also take a systematic approach to analyzing the sensitivity of our main estimates to using alternative samples, alternative instruments for causal mortality risk, alternative covariates, and alternative measures for the cost of homeownership. First, we consider four alternate

³³ These costs intentionally exclude the psychic cost of moving, which is thought to be important for residential choice, but is not a direct fiscal expense. The psychic cost of moving is also thought to be less important for within-state moves that most closely align with the residual variation in mortality risk and housing prices that identify the cost of mortality risk reduction in our econometric model.

³⁴ One could follow Bieri et al. (2023) to annualize this cost of moving over an individual’s remaining life years using a discount rate of 3% and age-specific life expectancies from the US Social Security Administration Actuarial Life Tables (US Social Security Administration, 2022). For an individual aged 77, this would yield an annualized cost of approximately \$1700, which is of similar magnitude to our estimated MCMRR (\$1474) of reducing mortality risk by 0.1 pp for the same age.

³⁵ In an alternative specification, we modified our dependent variable for homeowners to be the market value of their house plus the net present discounted value of expected lifetime property taxes using a 3% discount rate and the property-tax calculation method described in Bishop et al. (2023a). This changed our main results by less than 1%.

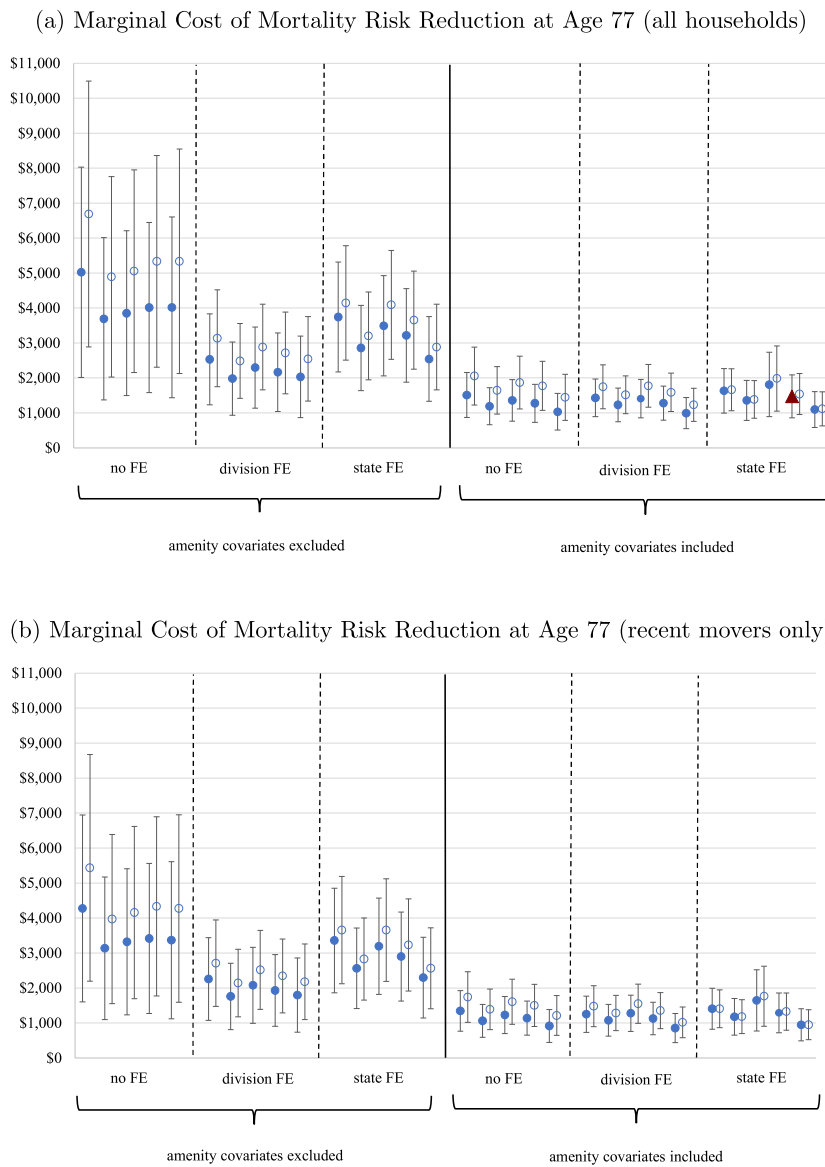


Fig. 3. Marginal cost of mortality risk reduction, sensitivity analysis. *Notes:* Each dot represents one of 120 estimates (varying over five different instruments, three different treatments of fixed effects, two different treatments of amenity covariates, and two different estimation samples) for the age-77 MCMRR, and is shown with its 95% confidence interval. Shaded dots follow our main specification in using owner-occupied housing prices. Unshaded dots use user-cost-of-housing prices that account for property taxes. The triangle denotes our main specification in column (4) of Table 1.

ways of defining the empirical mortality instrument: (i) CDC mortality at age 65, (ii) CDC mortality at age 75, (iii) CDC mortality at age 85+, and (iv) a weighted-average measure of empirical mortality from the same CMS sample that we use to estimate causal mortality following the same procedure as for our preferred instrument. The first three instruments use less information than our preferred instrument, but they avoid the need to aggregate over multiple ages. While the final instrument may embed measurement error due to working with a 20% sample, it has the advantage of using the same data to measure both causal and empirical mortality. We use each of these four instruments, plus our preferred instrument, to estimate models that differ in the following characteristics: (a) whether the model includes amenity covariates, (b) whether the model includes state dummies, Census division dummies, or no geographic dummies, (c) whether the model uses the full housing sample or just houses that changed occupant within the previous five years, and (d) whether owner-occupied house prices are adjusted by the user-cost rate to account for property taxes. With five alternative instruments, two samples, two sets of covariates,

three ways of modeling spatial dummies, and two ways of measuring the cost of homeownership, we estimate 120 different models.

Fig. 3 shows the sensitivity of our main estimates for the MCMRR at age 77 to the 120 alternative IV specifications. For reference, the result from our main specification in column (4) of Table 1 is shown as a triangle. Each dot in the figure reports an estimate and its 95% confidence interval from a specification that differs from our main specification in one or more of the following dimensions: (i) the instrumental variable; (ii) the estimation sample: all houses or only houses lived in by recent movers, (iii) the treatment of amenity covariates: included or excluded, and (iv) the treatment of fixed effects: none, Census divisions, or U.S. states. Within each delineated sub-panel the five pairs of dots correspond to the following instrumental variables ordered from left to right: (i) CDC empirical mortality for people age 65, (ii) CDC empirical mortality for people age 75, (iii) CDC empirical mortality for people aged 85 and above, (iv) our preferred CDC measure of empirical mortality, (v) an equivalent measure of empirical mortality calculated from the same CMS data that we use to estimate causal

mortality. Finally, within each dot pair, the shaded dot follows our main specification in using owner-occupied housing prices, whereas the unshaded dot adjusts housing prices by the user-cost rate to account for property taxes.

Comparing Figs. 3(a) to 3(b) shows that the results are robust to limiting the sample to people who moved to their current dwellings within the last five years and, therefore, may more accurately assess the market values of their houses. The three sub-panels in the left half of each sub-figure show that when amenity covariates are excluded from the model adding fixed effects for Census divisions or U.S. states reduces the MCMRR. Comparing the left half of each sub-figure to the right half shows that adding amenity covariates further reduces the estimated MCMRR. The right three sub-panels within each sub-figure show that when amenity covariates are included the MCMRR estimates are relatively robust to using alternative fixed effects and/or alternative empirical mortality instruments. Finally, comparing the shaded and unshaded dots within each pair shows that including property taxes as an additional cost of housing increases the MCMRR. However, the size of this increase is diminished by adding fixed effects or amenity covariates. When the model includes both state fixed effects and amenity covariates, as in our main specification, the effect of property taxes on the MCMRR ranges from less than 1% to approximately 9% depending on the instrument and sample.

For the 60 specifications that use amenity covariates, the MCMRR estimates fall within a fairly narrow range from \$855 to \$2132. At the top of this range, an MCMRR of \$2132 corresponds to a cost of \$0.8 million to avoid one premature statistical death at age 77.³⁶ Thus, the conclusion that our upper bound estimates are far below the cost of reducing mortality risk for workers and drivers over age 60 is robust to a wide range of econometric specifications.³⁷

6. Conclusion

This paper provides the first evidence that all-cause location-specific mortality risk is capitalized into housing prices. Specifically, we find the implicit housing cost is \$3453 (or less) for a 0.1 pp reduction in mortality risk among people in their late 60s and \$629 (or less) among people in their late 80s. Rescaling these household-level measures to individual measures yields results that are less than one-fifth of the conventional labor-market estimates for the cost of reducing mortality risk among workers in their early 60s and less than one-tenth of the cost of reducing mortality risk among drivers over age 60. By contrast, our results are similar to estimates for the medical cost of reducing mortality risk among the general population of Americans over age 65.

One hypothesis for the discrepancy in estimates across these markets is that people may be better informed about job-related and auto-related mortality risks than about location-based and health-based mortality risks. Another hypothesis is that the discrepancy reflects life-cycle heterogeneity in the willingness to pay for mortality-risk reduction, along with differences in health and wealth across study populations. Testing these hypotheses would require knowledge about households' beliefs about future spatial variation in mortality risk (Bishop and Murphy, 2019) and how these beliefs affect their migration decisions (Mathes, 2021). Developing this knowledge is an important area for further research.

³⁶ The same specification implies a cost of approximately \$2 million to avoid one premature statistical death at age 67.

³⁷ If we drop amenity covariates entirely, in order to force our MCMRR measure to absorb their effects on mortality risk along with their correlated effects on the quality of life, then the largest "extreme upper bound" estimate in the left half of Fig. 3(a) implies a cost of avoiding a premature statistical death at age 67 that is \$5.2 million, approximately half the size of labor-market and auto-market estimates for people over age 60.

Appendix. Supplemental online appendix

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jue.2023.103627>.

References

- Albouy, D., 2009. The unequal geographic burden of federal taxation. *J. Polit. Econ.* 117 (4), 635–667.
- Albouy, D., Graf, W., Kellogg, R., Wolff, H., 2016. Climate amenities, climate change, and American quality of life. *J. Assoc. Environ. Resour. Econ.* 3 (1), 205–246.
- Aldy, J.E., Viscusi, W.K., 2008. Adjusting the value of a statistical life for age and cohort effects. *Rev. Econ. Stat.* 90 (3), 573–581.
- Aliprantis, D., Richter, F.G.-C., 2020. Evidence of neighborhood effects from moving to opportunity: LATES of neighborhood quality. *Rev. Econ. Stat.* 102 (4), 633–647.
- Altonji, J.G., Conley, T., Elder, T.E., Taber, C.R., 2015. Methods for using selection on observed variables to address selection on unobserved variables.
- Altonji, J.G., Elder, T.E., Taber, C.R., 2005. Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *J. Political Econ.* 113 (1), 151–184.
- Angrist, J.D., Pischke, J.-S., 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Ansari, T., 2022. Where are people living the longest? See where your state ranks in life expectancy. *Wall Street J.* August 23.
- Bajari, P., Kahn, M.E., 2005. Estimating housing demand with an application to explaining racial segregation in cities. *J. Bus. Econom. Statist.* 23 (1), 20–33.
- Banzhaf, H.S., 2022. The value of statistical life: A meta-analysis of meta-analyses. *J. Benefit-Cost Anal.* 13 (2), 182–197.
- Banzhaf, S., Ma, L., Timmins, C., 2019. Environmental justice: The economics of race, place, and pollution. *J. Econ. Perspect.* 33 (1), 185–208.
- Banzhaf, H.S., Smith, V.K., 2007. Meta-analysis in model implementation: Choice sets and the valuation of air quality improvements. *J. Appl. Econometrics* 22, 1013–1031.
- Banzhaf, H.S., Walsh, R.P., 2008. Do people vote with their feet? An empirical test of tiebout. *Amer. Econ. Rev.* 98 (3), 843–863.
- Barreca, A., Clay, K., Deschênes, O., Greenstone, M., Shapiro, J.S., 2015. Convergence in adaptation to climate change: Evidence from high temperatures and mortality, 1900–2004. *Amer. Econ. Rev.* 105 (5), 247–251.
- Bayer, P., Ferreira, F., McMillan, R., 2007. A unified framework for measuring preferences for schools and neighborhoods. *J. Polit. Econ.* 115 (4), 588–638.
- Bayer, P., Keohane, N., Timmins, C., 2009. Migration and hedonic valuation: The case of air quality. *J. Environ. Econ. Manag.* 58 (1), 1–14.
- Bayer, P., Ross, S.L., Topa, G., 2008. Place of work and place of residence: Informal hiring networks and labor market outcomes. *J. Polit. Econ.* 116 (6), 1150–1196.
- Besley, T., Mueller, H., 2012. Estimating the peace dividend: The impact of violence on house prices in Northern Ireland. *Amer. Econ. Rev.* 102 (2), 810–833.
- Bieri, D.S., Kuminoff, N.V., Pope, J.C., 2023. National expenditures on local amenities. *J. Environ. Econ. Manag.* 117, 102717.
- Billings, S.B., Schnepel, K.T., 2017. The value of a healthy home: Lead paint remediation and housing values. *J. Public Econ.* 153, 69–81.
- Bishop, K.C., Dowling, J., Kuminoff, N.V., Murphy, A.D., 2023a. Tax policy and the heterogeneous costs of homeownership. NBER Working paper 31824.
- Bishop, K.C., Ketcham, J.D., Kuminoff, N.V., 2023b. Hazed and confused: The effect of air pollution on dementia. *Rev. Econom. Stud.* 90 (5), 2188–2214.
- Bishop, K.C., Murphy, A.D., 2019. Valuing time-varying attributes using the hedonic model: When is a dynamic approach necessary? *Rev. Econ. Stat.* 101 (1), 134–145.
- Blomquist, G.C., Berger, M.C., Hoehn, J.P., 1988. New estimates of quality of life in urban areas. *Am. Econ. Rev.* 89–107.
- Caetano, G., Macartney, H., 2021. What determines school segregation? The crucial role of neighborhood factors. *J. Public Econ.* 194, 104335.
- Card, D., Rothstein, J., Moises, Y., 2022. Location, Location, Location. Working paper.
- Case, A., Deaton, A., 2015. Rising morbidity and mortality among white and non-hispanic americans in the 21st century. *Proc. Natl. Acad. Sci.* 112 (49), 15078–15803.
- Chetty, R., Hendren, N., 2018a. The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects. *Q. J. Econ.* 133 (3), 1107–1162.
- Chetty, R., Hendren, N., 2018b. The impacts of neighborhoods on intergenerational mobility II: County-level estimates. *Q. J. Econ.* 133 (3), 1163–1228.
- Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., Cutler, D., 2016. The association between income and life expectancy in the united states, 2001–2014. *JAMA* 315 (16), 1750–1766.
- Christensen, P., Keiser, D., Lade, G., 2023. Economic effects of environmental crises: Evidence from Flint, Michigan. *Am. Econ. J.: Econ. Policy* 15 (1), 196–232.
- Cinelli, C., Forney, A., Pearl, J., 2022. A crash course in good and bad controls. *Sociol. Methods Res.*
- Costa, D.L., Kahn, M.E., 2004. Changes in the value of life, 1940–1980. *J. Risk Uncertain.* 29 (2), 159–180.
- Couillard, B.K., Foote, C.L., Gandhi, K., Meara, E., Skinner, J., 2021. Rising geographic disparities in U.S. mortality. *J. Econ. Perspect.* 35 (4), 123–146.

- Cropper, M., Hammitt, J.K., Robinson, L.A., 2011. Valuing mortality risk reductions: progress and challenges. *Annu. Rev. Resour. Econ.* 3 (1), 313–336.
- Currie, J., Schwandt, H., 2016. Mortality inequality: the good news from a county-level approach. *J. Econ. Perspect.* 30 (2), 29–52.
- Cutler, D.M., Glaeser, E.L., 1997. Are ghettos good or bad? *Q. J. Econ.* 112 (3), 827–872.
- Davis, L.W., 2004. The effect of health risk on housing values: Evidence from a cancer cluster. *Amer. Econ. Rev.* 94 (5), 1693–1704.
- Deryugina, T., Heutel, G., Miller, N.H., Molitor, D., Reif, J., 2019. The mortality and medical costs of air pollution: Evidence from changes in wind direction. *Amer. Econ. Rev.* 109 (12), 4178–4219.
- Deryugina, T., Molitor, D., 2020. Does when you die depend on where you live? Evidence from Hurricane Katrina. *Amer. Econ. Rev.* 110 (11), 3602–3633.
- Deryugina, T., Molitor, D., 2021. The causal effects of place on health and longevity. *J. Econ. Perspect.* 35 (4), 147–170.
- Deschenes, O., 2014. Temperature, human health, and adaptation: A review of the empirical literature. *Energy Econ.* 46, 606–619.
- Deschenes, O., Greenstone, M., Shapiro, J.S., 2017. Defensive investments and the demand for air quality: Evidence from the Nox budget program. *Amer. Econ. Rev.* 107 (10), 2958–2989.
- Diamond, R., 2016. The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000. *Amer. Econ. Rev.* 106 (3), 479–524.
- Doyle, J.J., Graves, J.A., Gruber, J., Kleiner, S.A., 2015. Measuring returns to hospital care: Evidence from ambulance referral patterns. *J. Polit. Econ.* 123 (1), 170–214.
- Ducharme, J., Wolfson, E., 2019. Your ZIP code might determine how long you live—and the difference could be decades. *Time Mag.* June 17.
- Dwyer-Lindgren, L., Bertozzi-Villa, A., Stubbs, R.W., Morozoff, C., Mackenbach, J.P., van Lenthe, F.J., Mokdad, A.H., Murray, C.J., 2017. Inequalities in life expectancy among US counties, 1980 to 2014: Temporal trends and key drivers. *JAMA Internal Med.* 177 (7), 1003–1011.
- Evans, M.F., Taylor, L.O., 2020. Using revealed preference methods to estimate the value of reduced mortality risk: Best practice recommendations for the hedonic wage model. *Rev. Environ. Econ. Policy* 14 (2), 282–301.
- Ferrari, N., 2017. 50 Ways to live a longer, healthier life. *AARP Bull.* (3).
- Finkelstein, A., Gentzkow, M., Williams, H.L., 2021. Place-based drivers of mortality: Evidence from migration. *Amer. Econ. Rev.* 111 (8), 2697–2735.
- Graff Zivin, J., Neidell, M., 2013. Environment, health, and human capital. *J. Econ. Lit.* 51 (3), 689–730.
- Hall, R.E., Jones, C.I., 2007. The value of life and the rise in health spending. *Q. J. Econ.* 122 (1), 39–72.
- Harding, J.P., Rosenthal, S.S., Sirmans, S., 2003. Estimating bargaining power in the market for existing homes. *Rev. Econ. Stat.* 85 (1), 178–188.
- Hausman, C., Stolper, S., 2022. Inequality, information failures, and air pollution. *J. Environ. Econ. Manag.* 102552.
- Huh, J., Reif, J., 2017. Did medicare part D reduce mortality? *J. Health Econ.* 53, 17–37.
- Jia, N., Molloy, R., Smith, C.L., Wozniak, A., 2023. The economics of internal migration: Advances and policy questions. *J. Econ. Lit.* 61 (1), 144–180.
- Kahn, M.E., 2004. Domestic pollution havens: Evidence from cancer deaths in border counties. *J. Urban Econ.* 56 (1), 51–69.
- Ketcham, J., Kuminoff, N., Saha, N., 2022. Valuing Statistical Life Using Seniors Medical Spending. Working paper.
- Kling, J.R., Liebman, J.B., Katz, L.F., 2007. Experimental analysis of neighborhood effects. *Econometrica* 75 (1), 83–119.
- Kniesner, T.J., Viscusi, W.K., Woock, C., Ziliak, J.P., 2012. The value of a statistical life: Evidence from panel data. *Rev. Econ. Stat.* 94 (1), 74–87.
- Kumbhaker, S.C., Parmeter, C.F., 2010. Estimation of hedonic price functions with incomplete information. *Empir. Econ.* 39, 1–25.
- Kuminoff, N.V., Smith, V.K., Timmins, C., 2013. The new economics of equilibrium sorting and policy evaluation using housing markets. *J. Econ. Lit.* 51 (4), 1007–1062.
- Lee, J.M., Taylor, L.O., 2019. Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life. *Am. Econ. J.: Econ. Policy* 11 (4), 350–374.
- Leggett, C.G., 2002. Environmental valuation with imperfect information. *Environ. Resour. Econ.* 23, 343–355.
- Lin, R., Ma, L., Phan, T., 2021. Race and Environmental Worries. Federal Reserve Bank of Richmond Working Paper, #21-15.
- Ma, L., 2019. Learning in a hedonic framework: Valuing brownfield remediation. *Internat. Econom. Rev.* 60 (3), 1355–1387.
- Mathes, S., 2021. The Dynamics of Residential Sorting and Health: Implications of Climate Change in the U.S. Working Paper.
- Muehlenbachs, L., Spiller, B., Timmins, C., 2015. The housing market impacts of shale gas development. *Amer. Econ. Rev.* 105 (12), 3633–3659.
- Nevo, A., Rosen, A.M., 2012. Identification with imperfect instruments. *Rev. Econ. Stat.* 94 (3), 659–671.
- O'Brien, J., 2018. Age, autos, and the value of a statistical life. *J. Risk Uncertain.* 57, 51–791013–31.
- Oster, E., 2019. Unobservable selection and coefficient stability: Theory and evidence. *J. Bus. Econom. Statist.* 37 (2), 187–204.
- Pope, J.C., 2008a. Buyer information and the hedonic: the impact of a seller disclosure on the implicit price for airport noise. *J. Urban Econ.* 63, 498–516.
- Pope, J.C., 2008b. Do seller disclosures affect property values? Buyer information and the hedonic model. *Land Econom.* 84 (4), 551–572.
- Pope, J.C., 2008c. Fear of crime and housing prices: household reactions to sex offender registries. *J. Urban Econ.* 64, 601–614.
- Portney, P.R., 1981. Housing prices, health effects, and valuing reductions in risk of death. *J. Environ. Econ. Manag.* 8 (1), 72–82.
- Poterba, J.M., 1984. Tax subsidies to owner-occupied housing: An asset-market approach. *Q. J. Econ.* 99 (4), 729–752.
- Roback, J., 1982. Wages, rents, and the quality of life. *J. Polit. Econ.* 90 (6), 1257–1278.
- Rohlf, C., Sullivan, R., Kniesner, T., 2015. New estimates of the value of a statistical life using air bag regulations as a quasi-experiment. *Am. Econ. J.: Econ. Policy* 7 (1), 331–359.
- Schlenker, W., Walker, W.R., 2016. Airports, air pollution, and contemporaneous health. *Rev. Econom. Stud.* 83 (2), 768–809.
- Shan, H., 2010. Property taxes and elderly mobility. *J. Urban Econ.* 67, 194–205.
- Smith, V.K., Evans, M.F., Kim, H., Taylor, Jr., D.H., 2004a. Do the near-elderly value mortality risks differently? *Rev. Econ. Stat.* 86 (1), 423–429.
- Smith, V.K., Sieg, H., Banzhaf, H.S., Walsh, R., 2004b. Estimating the general equilibrium benefits of large changes in spatially delineated public goods. *J. Environ. Econ. Manag.* 47, 559–584.
- US Social Security Administration, 2022. Annual Statistical Supplement to the Social Security Bulletin, 2022. SSA Publication No. 13-11700.
- Viscusi, W.K., Aldy, J.E., 2003. The value of a statistical life: a critical review of market estimates throughout the world. *J. Risk Uncertain.* 27 (1), 5–76.
- Weintraub, K., 2014. CDC: Lifespan more to do with geography than genetics. *USA Today* May 1.