FEDERAL RESERVE BANK OF ST. LOUIS

REVIEW

THIRD QUARTER 2017
VOLUME 99 | NUMBER 3

REVIEW

Volume 99 • Number 3

Review
Review is published four times per year by the Research Division of the Federal Reserve Bank of St. Louis. Complimentary print subscriptions are available to U.S. addresses only. Full online access is available to all, free of charge.

Online Access to Current and Past Issues

The current issue and past issues dating back to 1967 may be accessed through our Research Division website: **http://research.stlouisfed.org/publications/review**. All nonproprietary and nonconfidential data and programs for the articles written by Federal Reserve Bank of St. Louis staff and published in *Review* also are available to our readers on this website.

Review articles published before 1967 may be accessed through our digital archive, FRASER: **http://fraser.stlouisfed.org/publication/?pid=820**.

Review is indexed in Fed in Print, the catalog of Federal Reserve publications (**http://www.fedinprint.org/**), and in IDEAS/RePEc, the free online bibliography hosted by the Research Division (**http://ideas.repec.org/**).

Authorship and Disclaimer

The majority of research published in *Review* is authored by economists on staff at the Federal Reserve Bank of St. Louis. Visiting scholars and others affiliated with the St. Louis Fed or the Federal Reserve System occasionally provide content as well. *Review* does not accept unsolicited manuscripts for publication.

The views expressed in *Review* are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Subscriptions and Alerts

Single-copy subscriptions (U.S. addresses only) are available free of charge. Subscribe here: **https://research.stlouisfed.org/publications/review/subscribe/**.

Our monthly email newsletter keeps you informed when new issues of *Review*, *Economic Synopses*, *The Regional Economist*, and other publications become available; it also alerts you to new or enhanced data and information services provided by the St. Louis Fed. Subscribe to the newsletter here: **http://research.stlouisfed.org/newsletter-subscribe.html**.

Copyright and Permissions

Economic Data

General economic data can be obtained through FRED® (Federal Reserve Economic Data), our free database with about 400,000 national, international, and regional data series, including data for our own Eighth Federal Reserve District. You may access FRED through our website: **https://fred.stlouisfed.org**.

ISSN 0014-9187

Why Are Life-Cycle Earnings Profiles Getting Flatter?

B. Ravikumar and *Guillaume Vandenbroucke*

The authors present a simple, two-period model of human capital accumulation on the job and through college attainment. They use a calibrated version of the model to explain the observed flattening of the life-cycle earnings profiles of two cohorts of workers. The model accounts for more than 55 percent of the observed flattening for high school-educated and for college-educated workers. Two channels generate the flattening in the model: selection (or higher college attainment) and a higher skill price for the more recent cohort. Absent selection, the model would have accounted for no flattening for high school-educated workers and about 23 percent of the observed flattening for college-educated workers. (JEL E20, I26, J24, J31)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2017, *99*(3), pp. 245-57.
https://doi.org/10.20955/r.2017.245-257

Life-cycle earnings profiles are becoming flatter. The earnings growth over the life cycle for workers in the 1980 cohort was substantially less than the earnings growth for workers in the 1940 cohort (Table 1). We divide a worker's life cycle into two periods: ages 20 to 39 and ages 40 to 59. We refer to workers aged 20-39 in year t as cohort t. We compute life-cycle earnings growth as the ratio of earnings between the first and second periods of working life. For the 1940 cohort, real earnings grew between the first and second periods by a factor of 2.22 for the high school-educated workers and by 2.52 for the college-educated workers.[1] For the 1980 cohort, real earnings grew by a factor of only 1.31 for the high school-educated workers and by only 1.84 for the college-educated workers.

Put differently, Table 1 documents that the life-cycle earnings profiles are becoming flatter. The earnings growth of high school-educated workers in the 1980 cohort is 59.3 percent of the earnings growth of high school educated workers in the 1940 cohort. That is, the life-cycle earnings profile was 40.7 percent flatter for the high school-educated workers in the recent cohort ($1-\frac{1.31}{2.22}=40.7$ percent). Similarly, the life-cycle earnings profile for college-educated workers in the 1980 cohort was 26.9 percent flatter than that for college-educated workers in

B. Ravikumar is a vice president and deputy director of research and Guillaume Vandenbroucke is a research officer and economist at the Federal Reserve Bank of St. Louis.

Table 1

Earnings Growth Data

	High school	College
1940 Cohort earnings growth	2.22	2.52
1980 Cohort earnings growth	1.31	1.84
% Flattening	40.7	26.9

SOURCE: Integrated Public Use Microdata Series (IPUMS) and authors' calculations. Synthetic cohort, employed white men. The 1940 cohort consists of workers aged 20-39 in 1940, and the 1980 cohort consists of workers aged 20-39 in 1980.

the 1940 cohort.[2] Even though our focus is on only two cohorts and two education groups, the flattening of life-cycle earnings profiles is a more pervasive phenomenon. For evidence of flattening in earnings profiles for other cohorts, demographic groups, education groups, and finer age groups, see Kong, Ravikumar, and Vandenbroucke (2016).

In this paper, we present a simple, two-period model that attempts to quantitatively explain the flattening of life-cycle profiles. We use the framework of Ben-Porath (1967), which has been used extensively in the life-cycle earnings literature; see, for example, Heckman, Lochner, and Taber (1998) and Huggett, Ventura, and Yaron (2011). Workers accumulate human capital on the job and increase their earnings over the life cycle. The on-the-job human capital accumulation technology exhibits diminishing returns. Workers are heterogeneous in their endowed ability to accumulate human capital, and those with higher ability accumulate more human capital. Workers are partitioned by ability into two groups: high school educated and college educated. Those below a threshold level of ability are labeled high school educated, and the rest are labeled college educated. High school-educated workers begin their work lives right away with an exogenously given human capital that is positively correlated with their ability. College-educated workers begin their work lives later and decide how much human capital they want at the start of their work lives by choosing an amount of (goods) expenditures while in college. We will show later that for college-educated workers this implies a positive correlation between ability and (endogenous) start-of-work-life human capital.

Two well-known implications of the Ben-Porath model are important for our cross-cohort results. First, diminishing returns imply that workers with higher initial human capital accumulate on-the-job human capital at a slower rate and therefore experience slower earnings growth. Second, ability has two opposing effects on earnings growth. On the one hand, higher ability implies faster on-the-job human capital accumulation and hence higher earnings growth. This is the direct effect of ability. On the other hand, higher-ability workers also have higher human capital at the start of their work lives and hence experience slower earnings growth. In our quantitative exercise, the direct effect of ability dominates.

There are two exogenous differences between cohorts in our model: the skill price (or the price per unit of human capital) and the fraction of workers who are college graduates. Skill price affects the start-of-work-life human capital of the college-educated workers; the fraction of college graduates affects the average ability of college-educated workers and high school-educated workers. Recent cohorts face a higher skill price, so the marginal return to human capital is higher. Hence college-educated workers in the recent cohort incur more expenditures and start their work lives with higher human capital. The Ben-Porath model then implies that the earnings growth for such workers would be less in the recent cohort. Recent cohorts also have a higher fraction of workers who are college graduates, so the average ability, conditional on education, is lower in the recent cohort. Given that the direct effect of ability dominates, the Ben-Porath model then implies that on-the-job human capital is accumulated at a slower rate in the recent cohort by both the high school and college educated and hence the earnings growth is less for both groups.

Quantitatively, our model implies 22.6 percent flattening of the earnings profiles for high school graduates and 15 percent for college graduates. Thus, our model accounts for more than 55 percent of the observed flattening for both high school graduates (22.6/40.7 = 55.5 percent) and college graduates. In our model, selection, or the ability threshold for college graduates (measured by the fraction of college graduates), plays an important quantitative role. If the fraction of college graduates in the 1980 cohort were the same as that in the 1940 cohort, then our model would imply no flattening for high school graduates and barely 6 percent flattening for college graduates, or less than 23 percent of the observed flattening.

Our paper is related to Kambourov and Manovskii (2009), Guvenen and Kuruscu (2010), Hendricks (2015), and Jeong, Kim, and Manovskii (2015), who also document the flattening of earnings profiles of successive cohorts of workers. Our evidence goes farther back, starting with the 1920 birth cohort. Their models involve demographic changes, changes in occupational mobility, or skill-biased technical change. We propose a different, simple model of the flattening that is also consistent with the increase in college attainment.

1 THE MODEL

In our model, individuals live for two periods, age 1 and age 2. They differ in their ability, a, which is endowed to them in age 1 and remains constant throughout their lives. Ability, a, is distributed in the population according to the distribution function $A(a)$.

Each worker is endowed with a high school education and human capital $h_1(a)$, which is an increasing function of ability. We partition, exogenously, the population of workers into two groups. One, which we refer to as high school educated, starts working at the beginning of the first period with human capital $h_1(a)$. The other group, which we refer to as college educated, spends a fraction s of the first period acquiring a college education. This group can choose to spend resources to enhance the quality of their college education. Specifically, the human capital of individuals in this group after college is produced via the technology $G(k, h_1(a), a)$, where $h_1(a)$ is their endowed initial human capital and k represents goods spending. We assume that $G_1(k, h_1(a), a)$ is increasing in a.

Workers can accumulate human capital on the job, in the spirit of Ben-Porath (1967). The technology for accumulating human capital on the job is $F(nh, a)$, where $n \in (0,1]$ is time spent in human capital accumulation and h is human capital at the start of work life. We assume that F exhibits diminishing returns to nh and that $F_1(nh, a)$ is increasing in a.

We assume that the skill price, w (i.e., the price per unit of human capital), grows exogenously at rate g. There are perfect credit markets where workers can borrow and lend freely at the interest rate r. On-the-job human capital depreciates at rate δ.

The optimization problem of a high school-educated worker is

$$J^{\mathrm{HS}}(a)=\max_{n} wh^{\mathrm{HS}}(1-n)+\frac{1+g}{1+r}w\left[(1-\delta)h^{\mathrm{HS}}+F\left(nh^{\mathrm{HS}},a\right)\right], \tag{1}$$

$$\text{s.t.} \quad h^{HS}=h_1(a). \tag{2}$$

The optimization problem of a college-educated worker is

$$J^{\mathrm{CO}}(a)=\max_{k}\left\{\max_{n} wh^{\mathrm{CO}}(1-s-n)+\frac{1+g}{1+r}w\left[(1-\delta)h^{\mathrm{CO}}+F\left(nh^{\mathrm{CO}},a\right)\right]\right\}-k, \tag{3}$$

$$\text{s.t.} \quad h^{\mathrm{CO}}=G\left(k,h_1(a),a\right). \tag{4}$$

Note that the inner maximization in problem (3) is similar to the maximization in problem (1). It represents the worker's choice to accumulate human capital on the job. There are two differences between these two problems. First, college-educated workers have already spent a fraction s of their first period of life in college. They have only $1 - s$ units of time left to work or learn on the job, while high school-educated workers have 1 unit of time. Second, the human capital at the start of work life is not the same for high school- and college-educated workers. The former start with their endowed human capital $h^{\mathrm{HS}} = h_1(a)$, while the latter start with h^{CO}, which is a choice variable and depends on the amount of expenditures in college. The outer maximization in problem (3) describes the choice of goods spending in college.

1.1 Decisions

The first-order condition for n is the same for high school- and college-educated workers:

$$wh=\frac{1+g}{1+r}whF_1(nh,a),\ h=h^{HS} \tag{5}$$

and

$$wh=\frac{1+g}{1+r}whF_1(nh,a),\ h=h^{CO}, \tag{6}$$

where in each equation h is the human capital at the start of work life, the left-hand side is the marginal cost of increasing n, and the right-hand side is the discounted marginal benefit.[3] The start-of-work-life human capital h^{HS} is exogenous, but h^{CO} is endogenous (see below).

Two points are worth noting here. First, the optimal value of n determined by equation (5) or (6) is decreasing in h: Higher h implies lower human capital accumulation on the job due to diminishing returns to nh. Thus, individuals with higher start-of-work-life human capital have less earnings growth. Second, the optimal n is increasing in ability since $F_1(nh, a)$ is increasing in a.

The start-of-work-life human capital for college-educated workers, h^{CO}, is determined by the optimal choice of expenditures, k, in college. The first-order condition for k is

$$1 = w\left(1 - s + \frac{1+g}{1+r}(1-\delta)\right)G_1\left(k, h_1(a), a\right), \tag{7}$$

where the left-hand side is the marginal cost of college spending and the right-hand side is the marginal benefit, both measured in goods. Equation (7) implies that the optimal goods spending in college is increasing in the skill price, w, and in ability, a. Thus, conditional on w, individuals with higher ability acquire more human capital in college (higher h^{CO}), and, conditional on ability, individuals facing a higher skill price also acquire more human capital in college.

2 MODEL MECHANICS

Two exogenous variables characterize a cohort in our model. One is the partition of workers between high school- and college-educated workers. To operationalize this partition, we assign a threshold value of ability, a^*, and assign education according to

$$\text{Education} = \begin{cases} \text{High school} & \text{if } a < a^* \\ \text{College} & \text{if } a \geq a^*. \end{cases}$$

The proportion of workers with a college education is given by $1 - A(a^*)$. The ability distribution is assumed to be constant across cohorts, so the cohort with more college-educated workers is characterized by a lower value of a^*.

The other exogenous difference between cohorts is the level of the skill price, w, at the start of the work life. The growth rate of the skill price, g, is the same for each cohort. Each cohort is thus characterized by two numbers, a^* and w. Differences across cohorts in a^* and w generate differences in earnings growth of high school- and college-educated workers in each cohort.

2.1 Human Capital and Earnings Growth

Figure 1 presents a stylized description of the key outcomes of our model for either cohort. In the top panel, the line describes human capital at the start of the work life for high school- and college-educated workers. For high school-educated workers, the line plots the function $h_1(a)$ up to the threshold a^*; for college-educated workers, the line plots the function $G(k, h_1(a), a)$ at the optimal value of k from problem (3), above the threshold a^*. Note the discontinuity at a^*: The marginal worker would start work with more human capital after a college education.

In the bottom panel, the line describes earnings growth over the life cycle. Two points noted in Section 1.1 are worth repeating. First, conditional on education, earnings growth is increasing in ability. This is mainly due to the assumption that a positively affects human capital accumulation on the job. There is an opposing force, however. Human capital at the start of work life is increasing in a for both high school-educated and college-educated workers,

Figure 1

Human Capital and Earnings Growth in a Cohort

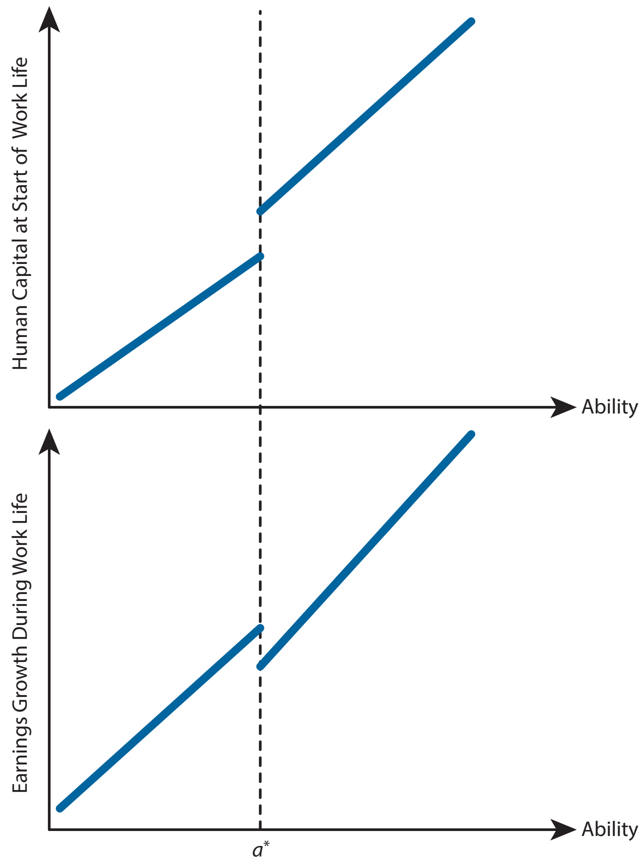

and the diminishing-returns property of F implies lower returns to human capital accumulation on the job. This effect may or may not be offset by the direct effect of ability. In Figure 1 we assume that the direct effect of ability offsets the return effect. This is also true in our quantitative exercise of Section 3.

Second, there is a discontinuity at a^*. The marginal worker would experience less earnings growth if he attended college than if he was high school educated only. This is because after a college education the marginal worker would start working with more human capital (top panel). This implies lower returns to human capital accumulation on the job and, in turn, less human capital accumulation and less earnings growth.

To compute the growth in average earnings of high school- and college-educated workers, we proceed as follows. Let $E_j^i(a)$ denote the earnings of an age j worker with ability a and education $i \in \{\mathrm{HS}, \mathrm{CO}\}$. Then,

$$E_j^{\mathrm{HS}} \equiv \int_0^{a^*} E_j^{\mathrm{i}}(a) A(da) \tag{8}$$

Figure 2

Cross-Cohort Differences: The Effect of Higher College Attainment, Holding the Skill Price Constant

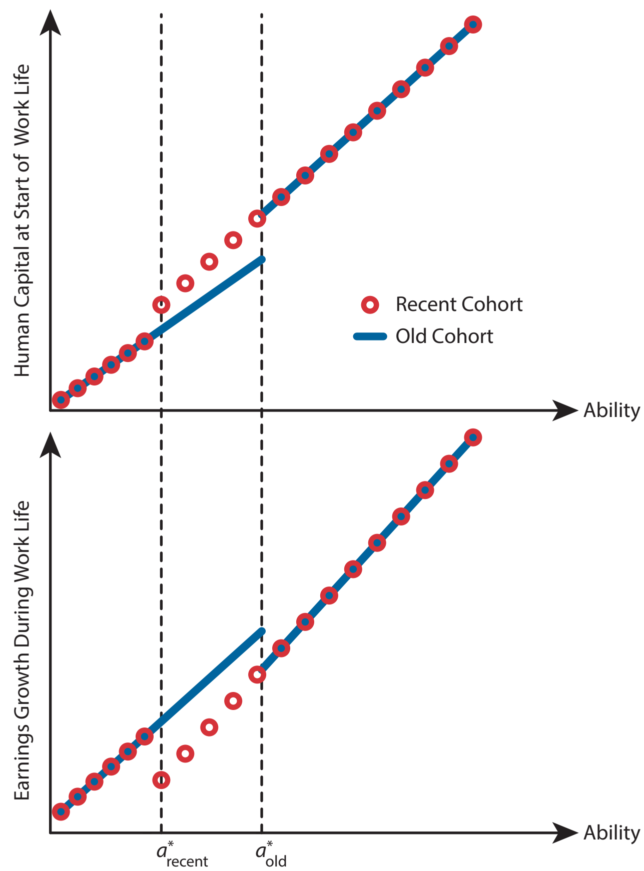

and

(9)
$$E_j^{\text{CO}} \equiv \int_{a^*} E_j^{\text{i}}(a)A(da).$$

In line with our definition of earnings growth in the Introduction, we measure earnings growth for workers with education i by E_2^i / E_1^i.

2.2 The Effect of College Attainment

Consider two cohorts, recent and old, that differ in the fraction of college-educated workers but face the same value for w. The recent cohort is characterized by a higher fraction of college-educated workers:

$$a^*_{\text{recent}} < a^*_{\text{old}}.$$

Figure 3

Cross-Cohort Differences: The Effect of Higher Skill Price, Holding College Attainment Constant

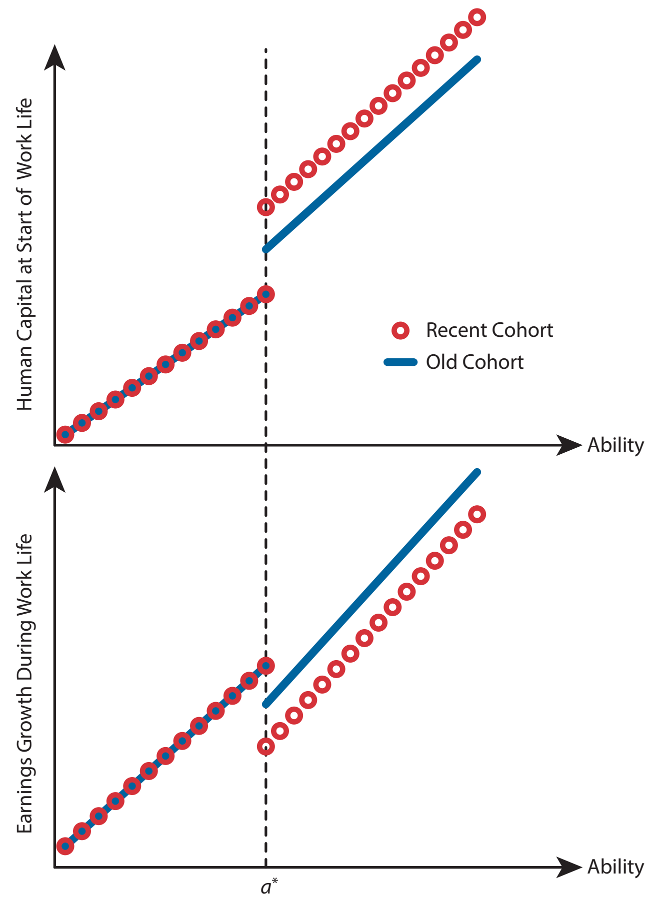

Figure 2 describes how human capital at the start of the work life and earnings growth differ for the two cohorts. The blue lines refer to the old cohort. The red circles refer to the recent cohort.

We adopt the following terminology: “always high school” is those workers with ability levels a such that $a < a^*_{\text{recent}}$ and “always-college” is those with ability levels a such that $a > a^*_{\text{old}}$. We use the term “switchers” to refer to those with ability $a \in [\, a^*_{\text{recent}} < a^*_{\text{old}} \,]$.

Note in Figure 2 that the human capital and earnings growth of always-high-school and always-college workers are the same for both cohorts. This is because these workers face exactly the same environment in the old and the recent cohorts. The switchers, however, acquire a college education in the recent cohort. As a result of a college education, they start their work lives with more human capital than their counterparts in the old cohort (top panel), which implies lower earnings growth over the life cycle (bottom panel).

It is clear from Figure 2 that the average earnings growth of both high school- and college-educated workers is lower in the recent cohort than in the old cohort. An increase in college attainment, therefore, implies a flattening of the average earnings profile of both high school- and college-educated workers. This is one channel for flattening in our model.

2.3 The Effect of Skill Price Growth

Consider now two cohorts that have the same college attainment but face different values for the skill price, w. With positive skill price growth, the recent cohort faces a higher value of w. Figure 3 describes the effects of this experiment.

Observe that human capital and earnings growth are the same for the high school-educated in both cohorts. This is because the environment for these workers is the same regardless of the level of w—this is clear from an inspection of problem (1). For college-educated workers, the level of the skill price in age 1 matters. A higher skill price implies goods are relatively cheaper, so the recent cohort would incur higher goods spending in college and, hence, finish college with more human capital (top panel). This implies, in turn, lower earnings growth for college-educated workers of the recent cohort (bottom panel).

It is thus clear from Figure 3 that the average earnings growth of high school-educated workers is the same in both cohorts and is not affected by the higher skill price. The average earnings growth of college-educated workers, however, is lower in the recent cohort since they start their work life with more human capital. This is the second channel for flattening in our model.

3 QUANTITATIVE IMPLICATIONS

We compare the life-cycle earnings profiles of workers in the 1940 and the 1980 cohorts. We measure flattening, as in Table 1, by the percentage change in the earnings growth over the life cycle; that is,

$$100\times\left(1-\frac{E_{2,1980}^{\mathrm{HS}}/E_{1,1980}^{\mathrm{HS}}}{E_{2,1940}^{\mathrm{HS}}/E_{1,1940}^{\mathrm{HS}}}\right)$$

for high school-educated workers and

$$100\times\left(1-\frac{E_{2,1980}^{\mathrm{CO}}/E_{1,1980}^{\mathrm{CO}}}{E_{2,1940}^{\mathrm{CO}}/E_{1,1940}^{\mathrm{CO}}}\right)$$

for college-educated workers.

3.1 Calibration

The functional forms for the human capital technologies are

$$h_1(a)=a, \tag{10}$$

$$G\left(k,h_1(a),a\right)=\left(z_G k\right)^{\eta}\left(a h_1(a)\right)^{1-\eta}, \tag{11}$$

$$F\left(nh,a\right)=z_F a\left(nh\right)^{\phi}. \tag{12}$$

The model period is 20 years. The rate of interest, r, is 5 percent per year. The rate of depreciation, δ, is 1 percent per year, similar to Huggett, Ventura, and Yaron (2006). The elasticity parameter in Equation (12) is $\phi = 0.75$; this is in the range of estimates reported by Browning, Hansen, and Heckman (1999). The growth rate of the skill price is 1.15 percent per year, the average growth rate used by Kong, Ravikumar, and Vandenbroucke (2016). The time spent in college is $s = 4/20$. The distribution of ability is log-normal:

$$\ln a \sim N\left(\mu,\sigma\right). \tag{13}$$

Finally, we choose a^*_{1940} and a^*_{1980} such that 30 percent of workers have a college education in the 1940 cohort and 50 percent in the 1980 cohort. These fractions correspond to college attainment for these cohorts in the U.S. data.

We adopt the normalization $\mu = 0$. The remaining parameters are z_F, z_G, σ, and η. To choose these parameters, we minimize the distance between moments implied by the model for the earnings of the 1940 cohort and their empirical counterparts. The moments we use are (i) earnings growth for the high school-educated workers; (ii) earnings growth for the college-educated workers; (iii) the coefficient of variation of earnings at age 2 for the high school-educated workers; and (iv) the coefficient of variation of earnings at age 2 for the college-educated workers. The empirical values of the first two moments are reported in Table 1. The empirical values for the last two moments are 0.5 and 0.6, respectively (see Kong, Ravikumar, and Vandenbroucke, 2016, Table 3).

3.2 Results

Table 2 reports the calibrated parameters. Table 3 reports our main results. Observe that the model implies earnings growth for the 1940 cohort that are the same as in the data. Although this is a target of the calibration, it should be noted that the earnings growth of a given cohort combines two effects. The first effect is the exogenous growth in the skill price: 1.15 percent per year. The second effect is the accumulation of human capital on the job. The annualized earnings growth rates in the model (and in the data) for the 1940 cohort are 4.3 percent and 5 percent per year for high school- and college-educated workers, respectively. Thus, the skill price growth is amplified nearly fourfold by the endogenous human capital accumulation in the model.

Comparing the 1940 and 1980 cohorts, the model delivers 22.6 percent flattening for the high school-educated workers; the corresponding figure in the data is 40.7 percent. Thus, the model accounts for 55.5 percent of the observed flattening for the high school-educated workers (22.6/40.7=55.5 percent). Similarly, the model delivers 15 percent flattening for college-educated workers. The figure in the data is 26.9 percent, so the model accounts for 55.7 percent of the observed flattening for college-educated workers.

Table 2

Calibration

A priori parameters		
Interest rate (annualized)	$r = 0.05$	
Depreciation (annualized)	$\delta = 0.01$	
Education	$s = 0.20$	
Ability distribution	$\mu = 0.00$	
On-the-job technology	$\phi = 0.75$	
Skill price growth (annualized)	$g = 0.0115$	
Calibrated parameters		
Ability distribution	$\sigma = 0.31$	
College technology	$z_G = 13.53$, $\eta = 0.42$	
On-the-job technology	$z_F = 2.51$	

Targeted moments (1940 cohort)	Data	Model
High school educated		
Earnings growth	2.22	2.22
Coefficient of variation of earnings	0.5	0.4
College educated		
Earnings growth	2.52	2.52
Coefficient of variation of earnings	0.6	0.6

Table 3

Baseline Results

	High school	College
1940 Cohort earnings growth	2.22	2.52
1980 Cohort earnings growth	1.72	2.14
% Flattening	22.6	15.0
% Flattening relative to data	55.5	55.7

Table 4

Counterfactual Experiment: No Change in College Attainment

	High school	College
1940 Cohort earnings growth	2.22	2.52
1980 Cohort earnings growth	2.22	2.36
% Flattening	0.0	6.1

Table 4 reports a decomposition: How much of the flattening implied by the model is due to selection? To answer this question, we hold college attainment to be the same in both cohorts. Specifically, the fraction of college-educated workers in the 1980 cohort is set at 30 percent, which is the observed fraction for the 1940 cohort. The skill price at the start of work life for the 1980 cohort is higher than that for the 1940 cohort, since the skill price grows annually at 1.15 percent. With college attainment the same for the two cohorts, the model implies *no* flattening for high school-educated workers (as noted in Figure 3) and only 6.1 percent flattening for college-educated workers. That is, without selection, the model accounts for none of the observed flattening for high school-educated workers and only 22.7 percent for college-educated workers.

4 CONCLUDING REMARKS

We document that life-cycle earnings profiles are getting flatter. For instance, the life-cycle earnings profile was about 41 percent flatter for high school-educated workers in the 1980 cohort relative to those in the 1940 cohort. And for college-educated workers in the 1980 cohort it was about 27 percent flatter relative to those in the 1940 cohort.

We develop a simple, two-period model of human capital accumulation on the job and through college attainment. We use a calibrated version of the model to explain the observed flattening of the life-cycle earnings profiles of two cohorts of workers. Our model accounts for more than 55 percent of the observed flattening for high school graduates and for college graduates. Two channels generate the flattening in our model: selection (or higher college attainment) and a higher skill price for the recent cohort. Absent selection, the model would have accounted for less than 23 percent of the observed flattening.

In our analysis, college enrollment is exogenous. Kong, Ravikumar, and Vandenbroucke (2016) present a multi-period life-cycle model where workers optimally choose whether to be college educated or not. Workers with ability above a (endogenously determined) threshold choose to become college educated and those below the threshold choose not to. Disciplining the model to match the time trend in college enrollment, they examine the flattening of life-cycle earnings profiles using cross-cohort differences in the skill price. ■

NOTES

[1] A period of working life is 20 years long in our calculations. The numbers in Table 1 correspond to a 4.3 percent annual rate of real earnings growth for high school-educated workers and 5 percent for college-educated workers.

[2] The numbers in Table 1 are based on synthetic cohorts of employed white men using Census data.

[3] A corner solution is possible. In this case, $n = 1$ for high school-educated workers and $n = 1 - s$ for college-educated workers. For our quantitative exercise in Section 3, the corner cases are not relevant.

REFERENCES

Ben-Porath, Yoram. "The Production of Human Capital and the Life-Cycle of Earnings." *Journal of Political Economy*, August 1967, *75*(4), pp. 352-65; https://doi.org/10.1086/259291.

Browning, Martin; Hansen, Lars Peter and Heckman, James J. "Micro Data and General Equilibrium Models." *Handbook of Macroeconomics*. Volume 1. Part A. Elsevier, 1999, pp. 543-633; https://doi.org/10.1016/S1574-0048(99)01011-3.

Guvenen, Fatih and Kuruscu, Burhanettin. "A Quantitative Analysis of the Evolution of the U.S. Wage Distribution, 1970-2000," in D. Acemoglu, K. Rogoff and M. Woodford, eds., *NBER Macroeconomics Annual 2009*. Volume 24. National Bureau of Economic Research, April 2010, pp. 227-76; http://www.nber.org/chapters/c11798.pdf.

Heckman, James J.; Lochner, Lance and Taber, Christopher. "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents." *Review of Economic Dynamics*, January 1998, *1*(1), pp.1-58; https://doi.org/10.1006/redy.1997.0008.

Hendricks, Lutz. "Accounting for Changing Returns to Experience." CFS Working Paper No. 558, Center for Financial Studies, October 2015; https://ssrn.com/abstract=2869179.

Huggett, Mark; Ventura, Gustavo and Yaron, Amir. "Human Capital and Earnings Distribution Dynamics." *Journal of Monetary Economics*, March 2006, *53*(2), pp. 265-90; https://doi.org/10.1016/j.jmoneco.2005.10.013.

Huggett, Mark; Ventura, Gustavo and Yaron, Amir. "Sources of Lifetime Inequality." *American Economic Review*, 2011, *101*(7), pp. 2923-54; https://doi.org/10.1257/aer.101.7.2923.

Jeong, Hyeok; Kim, Yong and Manovskii, Iourii. "The Price of Experience." *American Economic Association*, February 2015, *105*(2), pp. 784-815; https://www.aeaweb.org/articles?id=10.1257/aer.20080841.

Kambourov, Gueorgui and Manovskii, Iourii. "Accounting for the Changing Life-Cycle Profile of Earnings." Working paper, University of Pennsylvania, October 2009; http://economics.sas.upenn.edu/~manovski/papers/accounting_for_profiles.pdf.

Kong, Yu-Chien; Ravikumar, B. and Vandenbroucke, Guillaume. "Explaining Cross-Cohort Differences in Life Cycle Earnings." Working Paper No. 2015-035C, Federal Reserve Bank of Saint Louis, 2016; https://research.stlouisfed.org/wp/more/2015-035.

How Do Local Labor Markets Affect Retirement?

Leora Friedberg, Michael T. Owyang, Wei Sun, and Anthony Webb

Compared with prime-age workers, older workers face an easier path out of the labor force if they lose their jobs during a recession. However, premature job exits or earnings losses in the years leading up to retirement may be particularly devastating to retirement savings. The authors analyze the impact of recent business cycles on retirement using multifaceted job transitions of older workers. They focus on local labor markets because older workers are particularly unlikely to move for work. Surprisingly, the biggest effect of a higher local unemployment rate on older workers is to *raise* the propensity to stay in one's current job. Older workers have fewer voluntary transitions to new jobs when the unemployment rate rises, but they especially have fewer voluntary transitions out of the labor force. Thus, the direct effect of job loss in inducing earlier retirement during recessions is outweighed by retirement delays among those with jobs. (JEL J26, J23, J62)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2017, *99*(3), pp. 259-78.
https://doi.org/10.20955/r.2017.259-78

INTRODUCTION

Considering the vast attention given to studying the impact of business cycles on unemployment, much less emphasis has been given to business cycle effects on retirement. Compared with prime-age workers (ages 25-54), older workers face an easier path out of the labor force if they lose their jobs during a recession, especially if they can access Social Security and private pensions. However, premature job exits or earnings losses in the years leading up to retirement may be particularly devastating to efforts to save for old age.

This latter view is buttressed by a few studies on the impact of job loss on older workers. Older workers who experience layoffs suffer larger earnings losses upon finding new work

Leora Friedberg is an associate professor of economics at the University of Virginia. Michael T. Owyang is an assistant vice president and economist at the Federal Reserve Bank of St. Louis. Wei Sun is an associate professor of economics in the Hanqing Advanced Institute of Economics and Finance and the School of Finance at Renmin University of China. Anthony Webb is research director of the Retirement Equity Lab at the New School's Schwartz Center for Economic Policy Analysis. The authors thank Natalia Kolesnikova, Alicia H. Munnell, and Till von Wachter for helpful comments and Kristie Engemann, Shoghik Hovnanisyan, Dan Muldoon, and Hannah G. Shell for excellent research assistance. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium.

compared with younger workers or else they retire earlier than their peers on average (Chan and Stevens, 2001, 2004, and Couch and Plazcek, 2010) with fewer years to build up their retirement wealth. This finding raises concerns that the Great Recession may have been particularly devastating to workers near retirement age. Yet, the increase in the unemployment rate among workers 55 years of age and older, which rose from a low of 2.9 percent in early 2006 to 7.1 percent in late 2009, was considerably smaller than the increase in unemployment among prime-age workers.

While a handful of recent studies have increased our understanding of the impact of business cycles on retirement, they have not fully incorporated insights from the extensive retirement literature. We make two contributions in this regard. First, we measure business cycles at the local level, focusing on metropolitan statistical areas (MSAs). Some studies have used national, state-level, and/or industry-level unemployment rates (Coile and Levine, 2007, 2011, and von Wachter, 2007) largely because of data limitations, but few studies have focused on the MSA level, which may be most relevant because older workers are particularly unlikely to move to seek a new job. Second, rather than treating retirement as a binary outcome, we consider several possible transitions—exiting the current job on a voluntary or involuntary basis, obtaining a new job that involves part-time or full-time hours, or leaving work entirely. A binary retirement variable does not fit the multifaceted retirement process well (Ruhm, 1990, and Gustman and Steinmeier, 1995). Some workers retire abruptly and others gradually phase into retirement by reducing work hours, often with an accompanying change in employer. Gradual retirement may be particularly attractive during a recession to older workers who are pushed out of full-time jobs. Moreover, labor demand factors, though often ignored in the retirement literature, can also affect late-life work options. We use information on voluntary and involuntary job exits to help characterize demand effects on retirement decisions.

We model multifaceted job transitions in local labor markets using the Health and Retirement Study (HRS), which includes workers 50 years of age and older from 1992 on. The HRS is unique in offering both rich local identifiers on a restricted basis and a lengthy panel, so we are able to observe high-frequency transitions over a long period.[1] We use these data to estimate multinomial logits of job transitions. This approach is not a full dynamic programming exercise, which would incorporate even more complexity in the retirement process but would require modeling how workers form expectations about future unemployment rates as well as their own income, assets, and health. Yet, the multinomial logit captures the richness of observed retirement transitions while maintaining a parsimonious specification.

Surprisingly, we find that the biggest effect of a higher MSA unemployment rate on older workers is to *raise* the propensity to stay in one's current job. Thus, the direct effect of job loss in inducing earlier retirement during recessions is outweighed by retirement delays among those with jobs. Older workers have fewer voluntary transitions to new jobs when the unemployment rate rises, which is unsurprising, but they also have fewer voluntary transitions out of the labor force. This increased attachment to the current job may reflect pessimism about economic well-being during retirement and a reduced fluidity of the labor market, which would undermine the option value of returning to work at a later date (Maestas, 2010). The estimated effects are relatively large, as a 1-percentage-point increase in the local unemployment rate increases the likelihood of staying in the same job by 0.49 percentage points and reduces the likelihood of voluntary retirement by 0.34 percentage points. By these estimates, the main impact of the Great Recession on older workers was to reduce voluntary retirement by 1.58 percentage points; this result fits with observations in the popular press highlighting the relatively smaller declines in employment among older workers during the Great Recession compared with prime-age workers. As one might expect, a higher local unemployment rate also raises involuntary job exits but, interestingly, those involuntary exits are equally likely to lead to a new job as to retirement within a one-year time frame.

We find similar overall impacts of local unemployment rates on older men and women, with men slightly more likely than women to stay in their current job when the unemployment rate rises. We also find mostly similar patterns for workers of different skill levels, though unskilled workers have a smaller increase in their propensity to stay in their current job when the unemployment rate is high and a lower decline in their propensity to retire voluntarily; semiskilled workers are most affected by the local unemployment rate.

In sum, our study sheds light on the influence of business cycles on retirement timing. Meanwhile, the local business cycle literature has concentrated on identifying differences in business cycles across metropolitan areas (Owyang et al., 2008; Owyang, Piger, and Wall, 2013; Wall, 2013), but it has paid less attention to differential effects when groups of workers differ in their mobility or willingness to opt out of the labor force permanently (Topel, 1986).

BACKGROUND

Postponing retirement is frequently touted as a solution to growing concerns related to financial well-being in old age, including inadequate retirement saving, post-retirement gaps in health insurance coverage, and underfunding of Social Security and Medicare. In planning retirement, at least half of workers state a desire to undertake a gradual transition from a full-time career job into retirement (U.S. General Accounting Office, 2001, and Hutchens, 2007). Opportunities for part-time work may facilitate postponing retirement, yet gradual retirement frequently necessitates a change of employer. Therefore, the ability of employees to exit the labor force at an age and in a manner of their choosing depends on local labor market conditions. Those conditions will affect both involuntary exits from jobs and the opportunity after involuntary or voluntary exits to find bridge jobs that allow phased retirement.

Although labor economists usually focus on the unemployment rate as a key characteristic of local labor markets, studies of retirement have ignored local labor demand until recently.[2] While retirement models have grown extraordinarily complex, the complexity arises in modeling individual budget constraints and preferences, rather than local conditions. As an example of what can be learned by incorporating both concerns, Black, Kolesnikova, and Taylor (2014) find that variation in commuting time helps explain large differences in married women's labor force participation rates across locations—even for women with the same number of children and levels of education.

Recent studies of retirement that directly or indirectly consider local labor markets offer a few exceptions. Chan and Stevens (2001, 2004) set the stage for consideration of labor market conditions by highlighting the extent to which involuntary job loss among older workers in the HRS spurs early retirement. They find that the probability of re-employment following displacement declines precipitously with age. Although they do not directly examine the role of local market conditions, their findings suggest labor market conditions may have been overlooked in the retirement literature. Similar findings appear in Coile and Levine (2011) and Callaway (2015), while Ozturk and Gallo (2013) show that job loss by older workers is associated with lower subsequent wealth accumulation. Black and Liang (2005) study the impact on older workers of shocks to county-specific steel and coal production and shocks to city-specific manufacturing production. Their empirical approach emphasizes natural experiments rather than estimation of full retirement models, in part because their data from the U.S. Census and Social Security Administration (SSA) lack the rich set of covariates available in the HRS.

Some recent work suggests that state-level economic conditions influence retirement, which underlines the importance of moving the focus to local conditions. Coile and Levine (2007, 2010, 2011) use 30 years of data from the Current Population Survey (CPS) to estimate the effects of state-level unemployment rates, along with stock market and real estate price changes, on retirement. They find that labor force exits of older workers, especially for those with a high school diploma, rise when state-level unemployment rates rise. However, their analysis does not distinguish how workers flow into retirement. von Wachter (2007) analyzes the impact of employment rates, also using data from the CPS, and finds results similar to those of Coile and Levine: When state-level unemployment rates rise, the employment of older workers declines. Complementary work by Munnell et al. (2008) uses data from the CPS from 1977 to 2007 and from the HRS to examine the role of state-level conditions.[3] However, the CPS has only a subset of the covariates available in the HRS and a very short panel.

A few recent articles directly or indirectly consider local labor market effects. Hairault, Langot, and Zylberberg (2015) include local fixed effects in a regression examining transitions into self-reported retirement in the HRS from employment versus unemployment; however, they do not incorporate variation in local labor market conditions over time.[4] Goda, Shoven, and Slavov (2012) use the HRS and find that county-level unemployment rates do not affect future retirement expectations, but this focus on future expectations fails to capture immediate business cycle effects arising from job loss, for example. The study most similar to ours, Maestas, Mullen, and Powell (2013), uses the HRS, as we do, to focus on local labor markets, but they emphasize characteristics of a person's current job, while we emphasize characteristics of the transition path that a person follows out of that job. They analyze industry-specific and nonpecuniary job characteristics in detail but consider only unidimensional transitions from employment to either non-employment or self-reported retirement.[5]

Compared with the recent literature, we model job transitions as a multidimensional process, which turns out to be critical in distinguishing business cycle effects on voluntary versus involuntary transitions. We also measure labor market conditions at the local level to evaluate how business cycles influence retirement transitions.

EMPIRICAL STRATEGY

Our approach involves estimating a multinomial logit model explaining annual job transitions for aging workers in the HRS. The emphasis in the literature on the heterogeneity in retirement transitions explains our multichotomous approach (Ruhm, 1990, and Gustman and Steinmeier, 1995). This approach is richer than common specifications that pick a single binary definition of retirement (leaving a career job, describing oneself as retired, working zero hours, and so on).[6] This also allows us to consider both voluntary and involuntary job exits, where we view voluntary job exits as reflecting labor supply factors and involuntary job exits as reflecting labor demand. This distinction has been overlooked in much of the retirement literature but is critical when evaluating local employment conditions.

Thus, we seek to explain the probability of observing outcome $y_{ntk} = 1, 2, \ldots K$ for each individual n in each year t, where the $K = 5$ outcomes at the end of the year are as follows:

- staying in the beginning-of-the-year job,
- leaving that job involuntarily to another job,
- leaving that job voluntarily to another job,
- leaving that job involuntarily to retirement, or
- leaving that job voluntarily to retirement.[7]

Ignoring for now the possible correlation of the error term across observations for the same individual, we can write $y_{ntk} = y_{ik}$. The probability that a particular y_{ik} is observed, conditional on observables x_i, can be expressed as

$$\Pr\left[y_{ik} = j \middle| x_i\right] = \frac{\exp\left(x_i{}'\beta_j\right)}{1+\sum_{j=1}^{K}\exp\left(x_i{}'\beta_j\right)}. \tag{1}$$

The covariates x_i reflect individual-level factors that affect retirement. These factors capture substitution and income effects on labor supply as well as labor demand conditions. This specification yields a coefficient estimate for each covariate x_i that is specific to each outcome k; so, for example, the local unemployment rate is allowed to have different effects on the likelihood of each possible transition but one. As is usual in the multinomial formulation, those coefficients are identified for K–1 of the outcomes, relative to an arbitrarily chosen outcome as a base case.

Relative to the structural retirement literature (e.g., Rust and Phelan, 1997; Gustman and Steinmeier, 2005; French and Jones, 2011), we do not (i) specify underlying preferences, (ii) model features of job outcomes that are not chosen, (iii) capture the full dynamics in the evolution of retirement benefits, or (iv) (specific to this case) model expectations about the future unemployment rate. Accounting for these issues carefully would require making functional form assumptions that tend to have little clear empirical justification. To deal with retirement benefits, we control parsimoniously for public and private pension characteristics associated with the gains to delaying exit from the current job (Coile and Gruber, 2007, and Friedberg and Webb, 2005). We also control for other characteristics of the initial job and of the individual, as described in the next section. We also allow for arbitrary correlation of the error term for observations that occur for the same individuals over time.

DATA

The HRS is a detailed longitudinal survey of over 7,600 representative households with a member born between 1931 and 1941. The HRS cohort began in 1992 and participants are surveyed every two years. We use data from the first 10 waves through 2010.[8]

The HRS asks about the precise timing of a job exit. We also use information on the reasons for leaving the former employer; if someone reports that the reason was a layoff or business closure or that they were encouraged to leave, the transition is classified as involuntary.[9] The HRS reports data on the zip code where each individual was interviewed at each wave, also on a restricted basis; the latter data enable us to assign individuals to local labor markets. Lastly, the HRS provides enormous detail about covariates that help explain retirement and may be correlated with local factors, such as job characteristics, health, marital status, and assets.

We define the individual's location as the core-based statistical area (CBSA) in which he or she was interviewed.[10] The HRS is intended to be nationally representative, subject to oversampling of minorities and residents of Florida.[11] Most micropolitan statistical areas (μSAs) and some small MSAs contain only a handful of respondents, although these contribute to our analysis of the overall impact of unemployment on labor market behavior. A potential difficulty with analyses of the impact of local labor market conditions on retirement transitions is the treatment of individuals who move from one MSA to another. In practice, this is not a significant issue. Among the person-year observations in our sample, only 1,217, or 3.5 percent, changed MSAs between one birthday and the next (which is when many people embark on retirement). We retain them in the sample, using their MSA at the beginning of the year to characterize their labor market conditions.

Our key geographic variable is the MSA-specific unemployment rate. The unemployment rate reflects conditions that workers face in the current job, in other potential jobs, and in retirement. As noted in the literature, older workers who lose their jobs may choose early retirement, and this early retirement channel is likely to deepen during recessions. However, a worse labor market also reduces voluntary job exits at any age because finding a new job is more difficult. A poor labor market may further slow voluntary retirement flows because resources available for retirement seem less solid and the option to return to the labor force becomes more difficult to exercise. We use unemployment rates for the period 1990-2010 obtained from the U.S. Bureau of Labor Statistics. As Figure 1 shows, the unemployment rate varies considerably across business cycles. The median MSA unemployment rate dropped steadily through the 1990s, reaching a low of 4.1 percent in 2000. After that it increased until 2003, then dropped, and then jumped to a high of 9.3 percent in 2010. The difference between the 25th and 75th percentile values is typically around 4 percentage points, while the minimum MSA unemployment rate is around 2 percent in any given year and the maximum typically exceeds 20 percent. We also find that location and year do not explain all of the variation in the unemployment rate, with a 21.7 percent residual variation remaining.

Figure 1

MSA Unemployment Rates

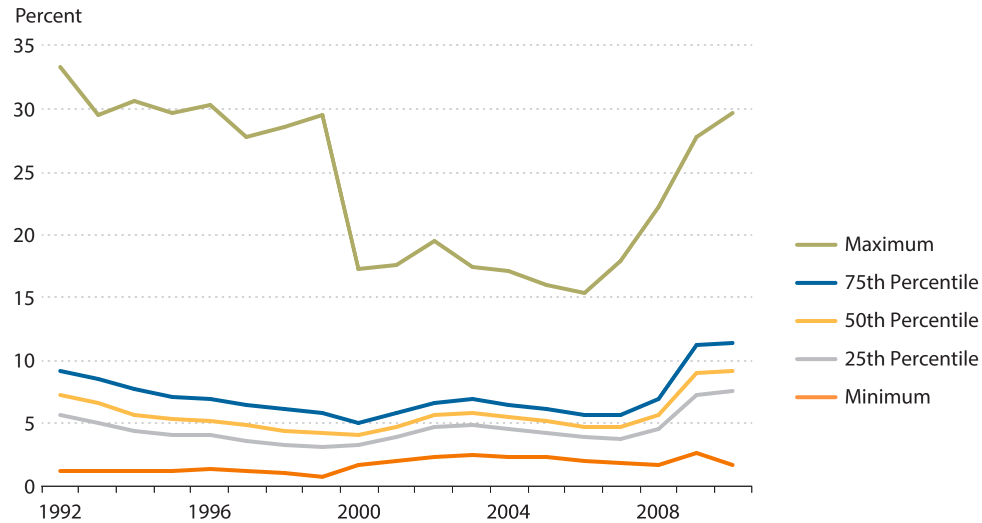

SOURCE: BLS Local Area Unemployment Statistics.

Our multinomial logit also includes demographic and socioeconomic controls that are known to explain retirement and that capture job characteristics related to the industrial composition of local labor markets. The controls include gender, marital status, race, education (three categories), self-reported health (five categories), single age dummies, financial wealth by quintile (which, though potentially endogenous, has little effect on other estimated coefficients when included), job tenure, plant size (six categories), industry (four categories), occupation (three categories),[12] whether the individual has responsibility for pay and promotion (a key indicator of management jobs), and union membership. We include information on employer-provided pensions. We use self-reported information on pension type (defined benefit, defined contribution, both, none).[13] Lastly, in other specifications we try to control for an individual's Social Security incentives.[14]

We select our sample as follows. Beginning with 12,652 individuals in the 1992 HRS, we keep individuals observed for at least one 12-month period starting at any age between 50 and 69, leaving 11,232 individuals. We drop those who were not working or self-employed in 1992, leaving 6,437.[15] Further, we drop those whose geographic identifiers are missing and who live outside an MSA or μSA, as their local unemployment rate cannot be obtained; we also exclude those for whom we cannot obtain financial or demographic data, leaving 5,387. We used the recall data on job transitions to convert the observations to 34,895 person-year observations. These person-year observations report the following employment transitions of workers from one birthday to the next: whether the person was working for the same employer and, if not, whether the person left voluntarily or involuntarily, stopped working, or took a job with another employer, and, if so, whether the new job involved full-time or part-time work hours (more or less than 30 weekly hours, respectively).[16]

For an idea of how those included in the sample transition out of their initial jobs and into retirement, note that between turning age 55 and turning age 56, 88.2 percent of the sample (defined as people who are in a job at the beginning of the period) stay in the same job. Among the rest, 2.6 percent lose their job involuntarily and take another job, 3.9 percent leave their job voluntarily and take another job, while 1.1 percent and 4.1 percent have the same types of exits, respectively, but retire. At age 60, almost the same percentage, 86.7 percent, stay in their jobs, while this declines to 85.0 percent at age 61 and 79.5 percent at age 62. Involuntary and voluntary job exits to another job decline gradually as the sample ages, while involuntary job exits to retirement remain roughly steady. Meanwhile, voluntary job exits to retirement rise to 6.3 percent at age 60, 9.2 percent at age 61, and 13.5 percent at age 62.

EMPIRICAL RESULTS

Interpretation of Multinomial Logit Results

Tables 1 through 5 report marginal effects and standard errors clustered at the person level obtained from weighted multinomial logit estimation. Our multinomial outcome variables involve birthday-to-birthday job transitions, as detailed below.[17] Each table reports the estimated marginal effects of the right-hand-side variables on the likelihood of one of the transitions occurring relative to the default state of staying in one's initial job. Each column of a table reports estimates for a different sample of interest, as detailed in the boxed insert.

The tables report the estimated effects of each covariate in the form of marginal effects. The marginal effect is a transformation of the estimated logit coefficient and captures the impact of a marginal change in a covariate on the likelihood of the occurrence of a particular job transition. As is common to multinomial models, the effect of covariates on the base outcome (staying in the current job) is not identified, as their coefficient estimates indicate the effect on the latent value of a particular outcome relative to the base outcome; marginal effects of these covariates on the base outcome are simply equal to 1 minus the sum of the marginal effects on the other outcomes.

Job transition outcome (reported by table)	Estimation sample (reported by column)
Stay in current job (Table 1)	Full sample (column (A))
Involuntary exit to a new job (Table 2)	Males (column (B))
Voluntary exit to a new job (Table 3)	Females (column (C))
Involuntary exit to retirement (Table 4)	Skilled occupations (column (D))
Voluntary exit to retirement (Table 5)	Semi-skilled occupations (column (E))
	Unskilled occupations (column (F))

Table 1

Multinomial Logit Marginal Effects: Outcome = Stay in Current Job

	(A)		(B)		(C)		(D)		(E)		(F)	
	Full sample		Males		Females		Skilled		Semi-skilled		Unskilled	
Variable	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Unemployment rate	0.0049	0.0010***	0.0057	0.0015***	0.0044	0.0013***	0.0048	0.0019**	0.0086	0.0019***	0.0031	0.0014**
Ages												
50-55	0.0466	0.0090***	0.0560	0.0148***	0.0469	0.0115***	0.0677	0.0164***	0.0488	0.0162***	0.0256	0.0148*
56-59	0.0352	0.0078***	0.0309	0.0123**	0.0421	0.0102***	0.0212	0.0136	0.0424	0.0139***	0.0423	0.0133***
60-61	–0.0059	0.0082	–0.0140	0.0123	0.0047	0.0112	–0.0094	0.0143	0.0225	0.0150	–0.0256	0.0136*
62	–0.0493	0.0093***	–0.0662	0.0144***	–0.0299	0.0123**	–0.0528	0.0160***	–0.0151	0.0177	–0.0757	0.0154***
63-64	–0.0255	0.0081***	–0.0300	0.0125**	–0.0188	0.0106*	–0.0385	0.0142***	0.0089	0.0148	–0.0428	0.0135***
65	–0.0444	0.0103***	–0.0528	0.0160***	–0.0370	0.0135***	–0.0600	0.0169***	–0.0260	0.0186	–0.0451	0.0183**
66	–0.0171	0.0114	–0.0285	0.0170*	–0.0082	0.0154	–0.0355	0.0191*	–0.0118	0.0207	–0.0063	0.0194
Male	0.0060	0.0049					0.0097	0.0081	-0.0039	0.0101	0.0089	0.0079
Education												
Less than high school	–0.0061	0.0059	0.0034	0.0092	–0.0124	0.0076	0.0039	0.0189	0.0144	0.0133	–0.0148	0.0077*
Some college	–0.0040	0.0050	–0.0064	0.0078	–0.0020	0.0066	0.0065	0.0089	–0.0046	0.0083	–0.0083	0.0091
Black	0.0088	0.0057	0.0037	0.0099	0.0120	0.0069*	–0.0342	0.0109***	0.0348	0.0148**	0.0194	0.0080**
Labor type												
Unskilled	–0.0009	0.0056	0.0096	0.0097	–0.0031	0.0068						
Skilled	0.0138	0.0060**	0.0265	0.0103**	0.0074	0.0075						
Industry												
Mining	–0.0080	0.0094	–0.0155	0.0114	0.0109	0.0215	0.0003	0.0191	–0.0151	0.0297	–0.0048	0.0126
Manufacturing	–0.0123	0.0057**	–0.0038	0.0079	–0.0228	0.0082***	–0.0024	0.0111	–0.0248	0.0109**	–0.0094	0.0090
Professional	0.0223	0.0055***	0.0218	0.0095**	0.0182	0.0066***	0.0328	0.0092***	0.0143	0.0098	0.0186	0.0095**
Married	–0.0011	0.0049	–0.0020	0.0098	–0.0024	0.0057	–0.0120	0.0086	–0.0010	0.0089	0.0064	0.0083
Union member	0.0043	0.0013***	0.0029	0.0020	0.0043	0.0018**	0.0022	0.0024	0.0072	0.0027***	0.0030	0.0021
Employee decides promotion	0.0054	0.0014***	0.0011	0.0019	0.0092	0.0019***	0.0027	0.0019	0.0071	0.0026***	0.0069	0.0028**
Size of work location												
< 5 employees	0.0067	0.0100	0.0029	0.0142	0.0099	0.0141	0.0217	0.0164	0.0225	0.0208	–0.0178	0.0172
5 to 14 employees	0.0030	0.0128	-0.0096	0.0199	0.0176	0.0172	–0.0125	0.0226	0.0045	0.0247	0.0254	0.0217
15 to 24 employees	0.0013	0.0121	0.0049	0.0168	–0.0003	0.0178	–0.0027	0.0207	0.0067	0.0258	0.0002	0.0189
25 to 99 employees	0.0151	0.0069**	0.0117	0.0113	0.0190	0.0088**	0.0060	0.0116	0.0223	0.0126*	0.0153	0.0119
100 to 499 employees	0.0026	0.0049	–0.0033	0.0074	0.0077	0.0066	0.0021	0.0082	0.0059	0.0093	–0.0028	0.0080
Financial wealth												
1st quintile	0.0146	0.0065**	0.0126	0.0098	0.0176	0.0086**	0.0318	0.0132**	0.0257	0.0119**	0.0026	0.0100
2nd quintile	–0.0018	0.0064	–0.0103	0.0096	0.0054	0.0086	–0.0026	0.0118	0.0061	0.0120	–0.0068	0.0103
3rd quintile	–0.0082	0.0062	–0.0265	0.0090***	0.0081	0.0085	–0.0108	0.0100	0.0028	0.0109	–0.0165	0.0113
4th quintile	–0.0015	0.0066	–0.0085	0.0100	0.0067	0.0089	–0.0044	0.0099	0.0160	0.0122	–0.0156	0.0127
Health												
Excellent	0.0204	0.0062***	0.0183	0.0092**	0.0217	0.0083***	0.0070	0.0095	0.0377	0.0117***	0.0192	0.0112*
Very good	0.0099	0.0049**	0.0129	0.0075*	0.0080	0.0065	0.0050	0.0084	0.0138	0.0092	0.0118	0.0082
Fair	–0.0247	0.0065***	–0.0210	0.0098**	–0.0257	0.0086***	–0.0176	0.0134	–0.0317	0.0128**	–0.0247	0.0094***
Poor	–0.0569	0.0147***	–0.0483	0.0211**	–0.0539	0.0220**	0.1569	0.0354***	–0.1016	0.0287***	–0.0506	0.0196***
Pension type												
DC	0.0436	0.0056***	0.0392	0.0084***	0.0451	0.0076***	0.0286	0.0093***	0.0416	0.0105***	0.0586	0.0094***
DB	0.0241	0.0061***	0.0160	0.0094*	0.0312	0.0082***	0.0129	0.0095	0.0270	0.0125**	0.0324	0.0105***
DB and DC	0.0067	0.0070	0.0060	0.0105	0.0052	0.0094	0.0040	0.0107	0.0166	0.0130	-0.0099	0.0127
Tenure	0.0014	0.0002***	0.0006	0.0003**	0.0024	0.0003***	0.0011	0.0004***	0.0023	0.0005***	0.0010	0.0004***
Marginal effect of unemployment rate at percentiles of the distribution												
Unemployment rate (p10)	0.0057	0.0012***	0.0068	0.0018***	0.0050	0.0015***	0.0058	0.0023**	0.0102	0.0025***	0.0038	0.0016**
Unemployment rate (p25)	0.0054	0.0011***	0.0064	0.0017***	0.0048	0.0014***	0.0054	0.0022**	0.0095	0.0022***	0.0035	0.0015**
Unemployment rate (p50)	0.0050	0.0010***	0.0059	0.0015***	0.0045	0.0013***	0.0049	0.0020**	0.0088	0.0020***	0.0032	0.0014**
Unemployment rate (p75)	0.0046	0.0009***	0.0053	0.0013***	0.0041	0.0012***	0.0043	0.0018**	0.0080	0.0017***	0.0029	0.0013**
Unemployment rate (p90)	0.0040	0.0008***	0.0046	0.0012***	0.0037	0.0011***	0.0036	0.0016**	0.0070	0.0014***	0.0025	0.0012**
N	34,895		15,327		19,568		10,884		9,711		14,297	

NOTE: Observations at the person-year level for a sample from the Health and Retirement Study, original HRS cohort, from 1992-2010. Coeff., coefficiant. S.E., standard error. The sample consists of those observed for at least two waves between ages 50-69 and initially employed or self-employed and who live in an MSA or μSA and for whom self-reports or imputations exist for all the variables in the regression. The table reports marginal effects estimated from a weighted multinomial logit (using survey weights to make the sample nationally representative) with five job transitions observed based on recall data from one birthday to the next (stay in current job, involuntary exit to new job, voluntary exit to new job, involuntary exit to retirement, voluntary exit to employment). The table also reports standard errors clustered at the person level, with statistical significance denoted by *** (1% level), ** (5% level), and * (10% level). The unemployment rate is measured at the MSA or μSA level. See the text for more information about control variables.

Table 2

Multinomial Logit Marginal Effects: Outcome = Involuntary Exit to a New Job

	(A) Full sample		(B) Males		(C) Females		(D) Skilled		(E) Semi-skilled		(F) Unskilled	
Variable	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Unemployment rate	0.0006	0.0003*	0.0004	0.0005	0.0007	0.0005	0.0011	0.0006*	−0.0001	0.0008	0.0007	0.0005
Ages												
50-55	0.0052	0.0036	0.0109	0.0057*	0.0014	0.0046	0.0028	0.0059	−0.0060	0.0070	0.0167	0.0061***
56-59	0.0036	0.0032	0.0074	0.0049	0.0014	0.0042	0.0013	0.0053	0.0008	0.0058	0.0086	0.0056
60-61	−0.0030	0.0037	−0.0015	0.0058	−0.0035	0.0048	0.0033	0.0061	−0.0111	0.0068	−0.0020	0.0063
62	−0.0087	0.0048*	−0.0115	0.0076	−0.0060	0.0061	−0.0043	0.0075	−0.0193	0.0097**	−0.0031	0.0079
63-64	−0.0058	0.0039	−0.0008	0.0060	−0.0098	0.0050**	0.0037	0.0060	−0.0225	0.0075***	−0.0016	0.0067
65	0.0020	0.0050	0.0012	0.0085	0.0037	0.0060	0.0049	0.0073	−0.0057	0.0091	0.0071	0.0092
66	−0.0108	0.0057*	−0.0072	0.0087	−0.0127	0.0077*	−0.0009	0.0086	−0.0200	0.0108*	−0.0113	0.0105
Male	0.0052	0.0022**					0.0053	0.0037	0.0093	0.0044**	0.0003	0.0036
Education												
Less than high school	−0.0018	0.0027	−0.0014	0.0043	−0.0025	0.0034	0.0000	0.0073	−0.0056	0.0067	−0.0003	0.0036
Some college	0.0050	0.0024**	0.0069	0.0041*	0.0033	0.0029	0.0004	0.0038	0.0084	0.0040**	0.0065	0.0044
Black	−0.0023	0.0028	−0.0055	0.0049	−0.0010	0.0033	0.0031	0.0055	−0.0134	0.0085	−0.0009	0.0036
Labor type												
Unskilled	−0.0043	0.0027	−0.0081	0.0046*	−0.0016	0.0030						
Skilled	−0.0040	0.0027	−0.0050	0.0047	−0.0048	0.0033						
Industry												
Mining	0.0050	0.0038	0.0074	0.0049	−0.0020	0.0082	0.0020	0.0072	0.0041	0.0109	0.0076	0.0052
Manufacturing	0.0020	0.0025	0.0015	0.0038	0.0028	0.0033	0.0038	0.0040	−0.0023	0.0052	0.0012	0.0045
Professional	−0.0113	0.0028***	−0.0128	0.0053**	−0.0094	0.0032***	−0.0117	0.0039***	−0.0069	0.0050	−0.0096	0.0049**
Married	−0.0024	0.0023	0.0063	0.0048	−0.0051	0.0027*	−0.0007	0.0040	−0.0058	0.0045	−0.0013	0.0037
Union member	0.0005	0.0007	0.0001	0.0011	0.0011	0.0009	0.0020	0.0014	0.0021	0.0018	−0.0009	0.0010
Employee decides promotion	−0.0003	0.0006	0.0004	0.0009	−0.0011	0.0009	−0.0010	0.0008	0.0005	0.0015	0.0009	0.0015
Size of work location												
< 5 employees	0.0020	0.0047	0.0076	0.0064	−0.0062	0.0076	−0.0050	0.0077	−0.0094	0.0125	0.0137	0.0069**
5 to 14 employees	0.0065	0.0060	0.0162	0.0089*	−0.0022	0.0083	−0.0104	0.0140	−0.0044	0.0138	0.0230	0.0085***
15 to 24 employees	−0.0002	0.0062	−0.0078	0.0106	0.0066	0.0074	−0.0008	0.0101	0.0076	0.0112	−0.0033	0.0110
25 to 99 employees	0.0009	0.0034	0.0093	0.0051*	−0.0055	0.0047	0.0054	0.0050	−0.0101	0.0070	0.0068	0.0061
100 to 499 employees	0.0079	0.0024***	0.0081	0.0036**	0.0078	0.0031**	0.0043	0.0034	0.0052	0.0047	0.0139	0.0043***
Financial wealth												
1st quintile	0.0008	0.0028	0.0041	0.0044	−0.0023	0.0037	−0.0007	0.0051	0.0013	0.0056	0.0004	0.0045
2nd quintile	0.0059	0.0029**	0.0082	0.0046*	0.0034	0.0037	0.0098	0.0045**	0.0065	0.0058	0.0032	0.0048
3rd quintile	0.0002	0.0029	0.0042	0.0044	−0.0029	0.0039	0.0013	0.0045	−0.0018	0.0057	0.0019	0.0051
4th quintile	0.0000	0.0031	0.0043	0.0049	−0.0034	0.0040	0.0006	0.0044	−0.0039	0.0060	0.0057	0.0061
Health												
Excellent	−0.0018	0.0028	0.0017	0.0043	−0.0046	0.0037	0.0050	0.0040	−0.0079	0.0057	−0.0037	0.0050
Very good	−0.0042	0.0023*	−0.0050	0.0039	−0.0032	0.0028	−0.0013	0.0039	−0.0066	0.0044	−0.0038	0.0040
Fair	−0.0041	0.0031	−0.0026	0.0048	−0.0050	0.0039	0.0010	0.0057	0.0028	0.0060	−0.0100	0.0047**
Poor	0.0002	0.0076	0.0077	0.0113	−0.0066	0.0089	0.0105	0.0121	0.0016	0.0146	−0.0020	0.0099
Pension type												
DC	−0.0075	0.0026***	−0.0064	0.0039	−0.0092	0.0035***	−0.0029	0.0034	−0.0121	0.0054**	−0.0095	0.0051*
DB	−0.0200	0.0036***	−0.0233	0.0056***	−0.0178	0.0047***	−0.0175	0.0053***	−0.0232	0.0082***	−0.0196	0.0055***
DB and DC	−0.0201	0.0037***	−0.0226	0.0057***	−0.0184	0.0050***	−0.0167	0.0048***	−0.0207	0.0071***	−0.0213	0.0074***
Tenure	−0.0008	0.0001***	−0.0007	0.0002***	−0.0009	0.0002***	−0.0007	0.0002***	−0.0012	0.0003***	−0.0006	0.0002***
Marginal effect of unemployment rate at percentiles of the distribution												
Unemployment rate (p10)	0.0006	0.0003*	0.0004	0.0005	0.0007	0.0004*	0.0010	0.0005**	0.0000	0.0008	0.0007	0.0005
Unemployment rate (p25)	0.0006	0.0003*	0.0004	0.0005	0.0007	0.0004	0.0010	0.0005**	0.0000	0.0008	0.0007	0.0005
Unemployment rate (p50)	0.0006	0.0003*	0.0004	0.0005	0.0007	0.0004	0.0011	0.0006*	0.0000	0.0008	0.0007	0.0005
Unemployment rate (p75)	0.0006	0.0004	0.0004	0.0005	0.0007	0.0005	0.0011	0.0006*	−0.0001	0.0008	0.0007	0.0005
Unemployment rate (p90)	0.0006	0.0004	0.0004	0.0006	0.0007	0.0005	0.0012	0.0007*	−0.0001	0.0008	0.0007	0.0006
N	34,895		15,327		19,568		10,884		9,711		14,297	

NOTE: Observations at the person-year level for a sample from the Health and Retirement Study, original HRS cohort, from 1992-2010. Coeff., coefficiant. S.E., standard error. The sample consists of those observed for at least two waves between ages 50-69 and initially employed or self-employed and who live in an MSA or μSA and for whom self-reports or imputations exist for all the variables in the regression. The table reports marginal effects estimated from a weighted multinomial logit (using survey weights to make the sample nationally representative) with five job transitions observed based on recall data from one birthday to the next (stay in current job, involuntary exit to new job, voluntary exit to new job, involuntary exit to retirement, voluntary exit to employment). The table also reports standard errors clustered at the person level, with statistical significance denoted by *** (1% level), ** (5% level), and * (10% level). The unemployment rate is measured at the MSA or μSA level. See the text for more information about control variables.

Table 3

Multinomial Logit Marginal Effects: Outcome = Voluntary Exit to a New Job

	(A)		(B)		(C)		(D)		(E)		(F)	
	Full sample		Males		Females		Skilled		Semi-skilled		Unskilled	
Variable	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
Unemployment rate	–0.0027	0.0006***	–0.0034	0.0010***	–0.0022	0.0007***	–0.0027	0.0011**	–0.0027	0.0011**	–0.0027	0.0008***
Ages												
50-55	0.0185	0.0045***	0.0289	0.0073***	0.0106	0.0056*	0.0246	0.0084***	0.0180	0.0083**	0.0157	0.0071**
56-59	0.0154	0.0043***	0.0256	0.0068***	0.0070	0.0054	0.0190	0.0081**	0.0126	0.0078	0.0165	0.0065**
60-61	0.0118	0.0045***	0.0178	0.0071**	0.0069	0.0057	0.0155	0.0084*	0.0061	0.0083	0.0143	0.0069**
62	0.0188	0.0050***	0.0331	0.0077***	0.0064	0.0067	0.0265	0.0092***	0.0079	0.0094	0.0214	0.0079***
63-64	0.0105	0.0046**	0.0169	0.0073**	0.0055	0.0059	0.0175	0.0088**	0.0043	0.0088	0.0109	0.0072
65	0.0116	0.0057**	0.0200	0.0088**	0.0053	0.0074	0.0248	0.0102**	0.0118	0.0102	–0.0010	0.0102
66	0.0060	0.0063	0.0112	0.0099	0.0023	0.0083	0.0106	0.0114	0.0106	0.0108	–0.0018	0.0106
Male	0.0064	0.0024***					0.0037	0.0038	0.0102	0.0053*	0.0071	0.0037*
Education												
Less than high school	–0.0080	0.0033**	–0.0155	0.0051***	–0.0016	0.0043	–0.0137	0.0116	–0.0105	0.0078	–0.0061	0.0042
Some college	0.0060	0.0026**	0.0020	0.0042	0.0088	0.0034***	0.0003	0.0048	0.0051	0.0046	0.0094	0.0042**
Black	–0.0045	0.0031	–0.0034	0.0052	–0.0041	0.0037	0.0048	0.0055	–0.0052	0.0075	–0.0089	0.0043**
Labor type												
Unskilled	0.0030	0.0029	0.0003	0.0052	0.0012	0.0034						
Skilled	–0.0002	0.0032	–0.0109	0.0055**	0.0060	0.0038						
Industry												
Mining	–0.0022	0.0044	–0.0008	0.0055	0.0041	0.0102	–0.0115	0.0102	0.0219	0.0115*	–0.0043	0.0057
Manufacturing	–0.0050	0.0029*	–0.0051	0.0044	–0.0043	0.0042	–0.0070	0.0057	–0.0055	0.0058	–0.0047	0.0045
Professional	–0.0024	0.0026	0.0057	0.0044	–0.0057	0.0031*	–0.0012	0.0042	–0.0033	0.0053	–0.0033	0.0046
Married	–0.0008	0.0025	0.0078	0.0050	–0.0023	0.0029	–0.0005	0.0046	0.0009	0.0047	–0.0029	0.0037
Union member	0.0012	0.0007*	0.0009	0.0010	0.0020	0.0010*	0.0019	0.0012	–0.0005	0.0015	0.0020	0.0011*
Employee decides promotion	–0.0002	0.0007	–0.0009	0.0010	0.0006	0.0010	0.0005	0.0009	0.0000	0.0015	–0.0018	0.0012
Size of work location												
< 5 employees	0.0075	0.0047	0.0051	0.0069	0.0098	0.0063	–0.0004	0.0080	0.0059	0.0091	0.0138	0.0075*
5 to 14 employees	0.0081	0.0066	–0.0004	0.0125	0.0129	0.0076*	0.0148	0.0109	0.0177	0.0102*	–0.0095	0.0130
15 to 24 employees	0.0071	0.0057	0.0055	0.0081	0.0080	0.0079	0.0122	0.0090	–0.0068	0.0139	0.0107	0.0086
25 to 99 employees	0.0048	0.0036	0.0069	0.0057	0.0031	0.0047	0.0058	0.0061	0.0067	0.0063	0.0029	0.0062
100 to 499 employees	0.0085	0.0026***	0.0079	0.0039**	0.0088	0.0034***	0.0069	0.0042*	0.0082	0.0051	0.0107	0.0042**
Financial wealth												
1st quintile	0.0002	0.0031	–0.0019	0.0049	0.0021	0.0040	0.0016	0.0058	–0.0055	0.0059	0.0046	0.0051
2nd quintile	0.0038	0.0031	0.0004	0.0048	0.0059	0.0040	0.0058	0.0055	0.0023	0.0056	0.0044	0.0050
3rd quintile	0.0012	0.0030	0.0074	0.0045*	–0.0048	0.0041	–0.0004	0.0049	0.0010	0.0054	0.0032	0.0055
4th quintile	–0.0080	0.0033**	–0.0004	0.0050	–0.0146	0.0045***	–0.0072	0.0050	–0.0141	0.0067**	–0.0012	0.0061
Health												
Excellent	0.0046	0.0030	0.0017	0.0044	0.0078	0.0039**	0.0035	0.0044	0.0012	0.0065	0.0087	0.0049*
Very good	0.0008	0.0026	–0.0063	0.0040	0.0070	0.0033**	–0.0040	0.0046	0.0034	0.0051	0.0022	0.0040
Fair	0.0013	0.0034	0.0019	0.0052	0.0009	0.0045	–0.0081	0.0073	0.0039	0.0065	0.0042	0.0049
Poor	–0.0123	0.0103	–0.0057	0.0134	–0.0225	0.0171	–0.0080	0.0161	0.0030	0.0197	–0.0194	0.0154
Pension type												
DC	–0.0149	0.0028***	–0.0150	0.0042***	–0.0147	0.0039***	–0.0096	0.0045**	–0.0162	0.0053***	–0.0190	0.005***
DB	–0.0171	0.0033***	–0.0154	0.0047***	–0.0208	0.0049***	–0.0152	0.0053***	–0.0191	0.0066***	–0.0156	0.0055***
DB and DC	–0.0126	0.0037***	–0.0154	0.0054***	–0.0102	0.0052**	–0.0094	0.0055*	–0.0220	0.0080***	–0.0062	0.0068
Tenure	–0.0008	0.0001***	–0.0006	0.0002***	–0.0010	0.0002***	–0.0008	0.0002***	–0.0010	0.0003***	–0.0006	0.0002**
Marginal effect of unemployment rate at percentiles of the distribution												
Unemployment rate (p10)	–0.0032	0.0008***	–0.0041	0.0014***	–0.0025	0.001***	–0.0031	0.0015**	–0.0032	0.0016**	–0.0033	0.0012***
Unemployment rate (p25)	–0.0030	0.0007***	–0.0038	0.0012***	–0.0024	0.0009***	–0.0029	0.0014**	–0.0030	0.0014**	–0.0031	0.0010***
Unemployment rate (p50)	–0.0027	0.0006***	–0.0034	0.0010***	–0.0022	0.0007***	–0.0027	0.0011**	–0.0028	0.0011**	–0.0028	0.0009***
Unemployment rate (p75)	–0.0025	0.0005***	–0.0030	0.0008***	–0.0020	0.0006***	–0.0024	0.0009***	–0.0025	0.0009***	–0.0025	0.0007***
Unemployment rate (p90)	–0.0022	0.0004***	–0.0026	0.0005***	–0.0018	0.0005***	–0.0021	0.0007***	–0.0022	0.0007***	–0.0022	0.0005***
N	34,895		15,327		19,568		10,884		9,711		14,297	

NOTE: Observations at the person-year level for a sample from the Health and Retirement Study, original HRS cohort, from 1992-2010. Coeff., coefficiant. S.E., standard error. The sample consists of those observed for at least two waves between ages 50-69 and initially employed or self-employed and who live in an MSA or μSA and for whom self-reports or imputations exist for all the variables in the regression. The table reports marginal effects estimated from a weighted multinomial logit (using survey weights to make the sample nationally representative) with five job transitions observed based on recall data from one birthday to the next (stay in current job, involuntary exit to new job, voluntary exit to new job, involuntary exit to retirement, voluntary exit to employment). The table also reports standard errors clustered at the person level, with statistical significance denoted by *** (1% level), ** (5% level), and * (10% level). The unemployment rate is measured at the MSA or μSA level. See the text for more information about control variables.

Table 4

Multinomial Logit Marginal Effects: Outcome = Involuntary Exit to Retirement

	(A) Full sample		(B) Males		(C) Females		(D) Skilled		(E) Semi-skilled		(F) Unskilled	
Variable	**Coeff.**	**S.E.**	**Coeff.**	**S.E.**	**Coeff.**	**S.E.**	**Coeff.**	**S.E.**	**Coeff.**	**S.E.**	**Coeff.**	**S.E.**
Unemployment rate	0.0006	0.0002***	0.0008	0.0003**	0.0003	0.0003	0.0007	0.0005	0.0005	0.0006	0.0006	0.0003**
Ages												
50-55	–0.0018	0.0031	–0.0079	0.0055	0.0012	0.0041	–0.0047	0.0061	–0.0065	0.0057	0.0054	0.0051
56-59	–0.0006	0.0027	–0.0028	0.0038	0.0005	0.0038	0.0012	0.0041	–0.0049	0.0051	0.0021	0.0047
60-61	0.0044	0.0028	0.0064	0.0037*	0.0018	0.0041	0.0074	0.0042*	–0.0071	0.0057	0.0114	0.0049**
62	0.0069	0.0031**	0.0040	0.0045	0.0085	0.0044*	0.0066	0.0047	0.0017	0.0063	0.0129	0.0054**
63-64	0.0054	0.0028*	0.0064	0.0038*	0.0041	0.0042	0.0071	0.0043*	–0.0002	0.0055	0.0096	0.0050*
65	0.0053	0.0036	0.0030	0.0055	0.0067	0.0049	0.0004	0.0063	0.0010	0.0073	0.0129	0.0059**
66	0.0029	0.0037	0.0029	0.0051	0.0023	0.0054	0.0014	0.0060	–0.0020	0.0077	0.0097	0.0061
Male	–0.0011	0.0016					0.0001	0.0024	–0.0021	0.0036	–0.0016	0.0026
Education												
Less than high school	0.0027	0.0018	0.0034	0.0028	0.0016	0.0024	–0.0013	0.0053	–0.0042	0.0050	0.0068	0.0026***
Some college	0.0004	0.0017	0.0023	0.0026	–0.0012	0.0024	0.0005	0.0026	–0.0001	0.0030	0.0016	0.0033
Black	0.0009	0.0020	0.0038	0.0031	–0.0013	0.0025	0.0076	0.0029***	–0.0164	0.0073**	0.0020	0.0028
Labor type												
Unskilled	–0.0046	0.0018***	–0.0059	0.0029**	–0.0032	0.0023						
Skilled	–0.0022	0.0020	–0.0016	0.0033	–0.0032	0.0028						
Industry												
Mining	0.0050	0.0029*	0.0069	0.0036*	–0.0094	0.0084	0.0072	0.0059	–0.0084	0.0158	0.0044	0.0034
Manufacturing	0.0055	0.0018***	0.0039	0.0026	0.0077	0.0025***	0.0032	0.0030	0.0102	0.0037***	0.0043	0.0028
Professional	–0.0086	0.0021***	–0.0078	0.0038**	–0.0090	0.0025***	–0.0078	0.0028***	–0.0117	0.0041***	–0.0075	0.0036**
Married	–0.0023	0.0016	–0.0039	0.0028	–0.0018	0.0020	–0.0014	0.0025	–0.0005	0.0032	–0.0040	0.0027
Union member	0.0005	0.0005	0.0009	0.0007	0.0001	0.0006	0.0001	0.0008	0.0004	0.0010	0.0009	0.0007
Employee decides promotion	0.0016	0.0006***	0.0023	0.0008***	0.0008	0.0008	0.0013	0.0006**	0.0028	0.0014**	0.0011	0.0012
Size of work location												
< 5 employees	–0.0031	0.0035	–0.0030	0.0047	–0.0031	0.0053	–0.0027	0.0057	–0.0069	0.0075	0.0000	0.0053
5 to 14 employees	0.0015	0.0042	0.0017	0.0059	0.0011	0.0061	0.0073	0.0052	–0.0039	0.0087	–0.0005	0.0078
15 to 24 employees	0.0047	0.0036	0.0023	0.0052	0.0072	0.0050	0.0083	0.0052	–0.0036	0.0090	0.0068	0.0055
25 to 99 employees	0.0001	0.0023	–0.0090	0.0043**	0.0045	0.0029	0.0011	0.0036	–0.0025	0.0044	0.0010	0.0039
100 to 499 employees	–0.0003	0.0017	0.0005	0.0025	–0.0011	0.0024	0.0016	0.0026	–0.0070	0.0035**	0.0035	0.0029
Financial wealth												
1st quintile	–0.0010	0.0021	0.0025	0.0031	–0.0044	0.0029	–0.0032	0.0038	–0.0040	0.0042	0.0028	0.0037
2nd quintile	–0.0035	0.0023	–0.0010	0.0035	–0.0058	0.0030*	–0.0061	0.0040	–0.0060	0.0044	0.0000	0.0039
3rd quintile	0.0019	0.0022	0.0050	0.0031	–0.0007	0.0031	0.0015	0.0029	–0.0061	0.0045	0.0096	0.0040**
4th quintile	0.0003	0.0023	0.0001	0.0034	0.0000	0.0031	–0.0002	0.0030	–0.0043	0.0045	0.0053	0.0045
Health												
Excellent	–0.0064	0.0022***	–0.0077	0.0035**	–0.0053	0.0029*	–0.0053	0.0033	–0.0128	0.0047***	–0.0010	0.0035
Very good	–0.0029	0.0017*	–0.0036	0.0024	–0.0025	0.0023	0.0012	0.0025	–0.0055	0.0035	–0.0058	0.0030*
Fair	0.0025	0.0021	–0.0002	0.0036	0.0046	0.0027*	0.0066	0.0042	0.0021	0.0045	0.0011	0.0030
Poor	0.0090	0.0036**	0.0052	0.0054	0.0124	0.0048***	–0.1988	0.0219***	0.0150	0.0079*	0.0109	0.0047**
Pension type												
DC	–0.0066	0.0020***	–0.0053	0.0030*	–0.0077	0.0028***	–0.0016	0.0028	–0.0097	0.0045**	–0.0100	0.0037***
DB	–0.0026	0.0020	–0.0043	0.0031	–0.0006	0.0026	–0.0080	0.0036**	0.0063	0.0037*	–0.0053	0.0035
DB and DC	–0.0004	0.0022	0.0023	0.0033	–0.0028	0.0031	0.0009	0.0029	–0.0059	0.0049	0.0028	0.0041
Tenure	–0.0002	0.0001***	–0.0001	0.0001	–0.0003	0.0001**	–0.0001	0.0001	–0.0002	0.0002	–0.0002	0.0001*
Marginal effect of unemployment rate at percentiles of the distribution												
Unemployment rate (p10)	0.0006	0.0002***	0.0007	0.0002***	0.0003	0.0003	0.0006	0.0004	0.0005	0.0005	0.0006	0.0002**
Unemployment rate (p25)	0.0006	0.0002***	0.0007	0.0003***	0.0003	0.0003	0.0006	0.0004	0.0005	0.0005	0.0006	0.0003**
Unemployment rate (p50)	0.0006	0.0002***	0.0008	0.0003**	0.0003	0.0003	0.0007	0.0005	0.0005	0.0005	0.0006	0.0003**
Unemployment rate (p75)	0.0006	0.0002**	0.0008	0.0003**	0.0003	0.0003	0.0007	0.0006	0.0005	0.0006	0.0006	0.0003**
Unemployment rate (p90)	0.0006	0.0003**	0.0009	0.0004**	0.0003	0.0004	0.0008	0.0007	0.0005	0.0006	0.0007	0.0004*
N	34,895		15,327		19,568		10,884		9,711		14,297	

NOTE: Observations at the person-year level for a sample from the Health and Retirement Study, original HRS cohort, from 1992-2010. Coeff., coefficiant. S.E., standard error. The sample consists of those observed for at least two waves between ages 50-69 and initially employed or self-employed and who live in an MSA or μSA and for whom self-reports or imputations exist for all the variables in the regression. The table reports marginal effects estimated from a weighted multinomial logit (using survey weights to make the sample nationally representative) with five job transitions observed based on recall data from one birthday to the next (stay in current job, involuntary exit to new job, voluntary exit to new job, involuntary exit to retirement, voluntary exit to employment). The table also reports standard errors clustered at the person level, with statistical significance denoted by *** (1% level), ** (5% level), and * (10% level). The unemployment rate is measured at the MSA or μSA level. See the text for more information about control variables.

Table 5

Multinomial Logit Marginal Effects: Outcome = Voluntary Exit to Retirement

Variable	(A) Full sample Coeff.	S.E.	(B) Males Coeff.	S.E.	(C) Females Coeff.	S.E.	(D) Skilled Coeff.	S.E.	(E) Semi-skilled Coeff.	S.E.	(F) Unskilled Coeff.	S.E.
Unemployment rate	–0.0034	0.0007***	–0.0036	0.0011***	–0.0033	0.001***	–0.0039	0.0015**	–0.0063	0.0015***	–0.0017	0.0010*
Ages												
50-55	–0.0685	0.0069***	–0.0880	0.0112***	–0.0600	0.0089***	–0.0903	0.0124***	–0.0542	0.0122***	–0.0634	0.0112***
56-59	–0.0536	0.0057***	–0.0610	0.0087***	–0.0511	0.0075***	–0.0427	0.0093***	–0.0509	0.0102***	–0.0695	0.0098***
60-61	–0.0074	0.0055	–0.0087	0.0082	–0.0099	0.0075	–0.0168	0.0094*	–0.0104	0.0100	0.0020	0.0091
62	0.0323	0.0058***	0.0406	0.0083***	0.0210	0.0082**	0.0241	0.0101**	0.0248	0.0106**	0.0445	0.0096***
63-64	0.0155	0.0055***	0.0076	0.0082	0.0190	0.0074***	0.0102	0.0094	0.0095	0.0097	0.0239	0.0092***
65	0.0255	0.0067***	0.0287	0.0099***	0.0212	0.0092**	0.0298	0.0114***	0.0190	0.0120	0.0261	0.0115**
66	0.0190	0.0072***	0.0215	0.0104**	0.0163	0.0099*	0.0243	0.0120**	0.0232	0.0126*	0.0097	0.0125
Male	–0.0165	0.0036***					–0.0188	0.0059***	–0.0135	0.0071*	–0.0147	0.0058**
Education												
Less than high school	0.0132	0.0042***	0.0101	0.0060*	0.0149	0.0059**	0.0111	0.0141	0.0059	0.0094	0.0143	0.0055***
Some college	–0.0074	0.0036**	–0.0048	0.0053	–0.0089	0.0050*	–0.0077	0.0066	–0.0087	0.0058	–0.0091	0.0067
Black	–0.0028	0.0042	0.0014	0.0071	–0.0057	0.0053	0.0188	0.0081**	0.0002	0.0091	–0.0117	0.0060*
Labor type												
Unskilled	0.0068	0.0040*	0.0041	0.0067	0.0067	0.0052						
Skilled	–0.0074	0.0043*	–0.0090	0.0076	–0.0055	0.0054						
Industry												
Mining	0.0002	0.0070	0.0019	0.0080	–0.0037	0.0180	0.0019	0.0148	–0.0025	0.0221	–0.0030	0.0093
Manufacturing	0.0099	0.0042**	0.0035	0.0057	0.0166	0.0064***	0.0024	0.0087	0.0224	0.0075***	0.0085	0.0065
Professional	0.0001	0.0039	–0.0069	0.0066	0.0059	0.0049	–0.0121	0.0071*	0.0076	0.0066	0.0019	0.0067
Married	0.0066	0.0036*	–0.0082	0.0063	0.0117	0.0043***	0.0146	0.0063**	0.0064	0.0062	0.0018	0.0058
Union member	–0.0067	0.0009***	–0.0049	0.0013***	–0.0074	0.0011***	–0.0061	0.0015***	–0.0093	0.0015***	–0.0050	0.0014***
Employee decides promotion	–0.0064	0.0009***	–0.0030	0.0014**	–0.0096	0.0013***	–0.0035	0.0013***	–0.0104	0.0015***	–0.0071	0.0018***
Size of work location												
< 5 employees	–0.0131	0.0074*	–0.0125	0.0102	–0.0105	0.0105	–0.0136	0.0116	–0.0120	0.0143	–0.0097	0.0124
5 to 14 employees	–0.0191	0.0094**	–0.0079	0.0134	–0.0294	0.0129**	0.0008	0.0149	–0.0139	0.0179	–0.0384	0.0153**
15 to 24 employees	–0.0129	0.0088	–0.0049	0.0113	–0.0214	0.0135	–0.0170	0.0158	–0.0038	0.0171	–0.0145	0.0133
25 to 99 employees	–0.0209	0.0047***	–0.0189	0.0077**	–0.0210	0.0060***	–0.0183	0.0081**	–0.0164	0.0085*	–0.0261	0.0078***
100 to 499 employees	–0.0187	0.0033***	–0.0131	0.0050***	–0.0233	0.0045***	–0.0148	0.0055***	–0.0124	0.0061**	–0.0253	0.0055***
Financial wealth												
1st quintile	–0.0147	0.0049***	–0.0173	0.0076**	–0.0131	0.0065**	–0.0295	0.0109***	–0.0175	0.0087**	–0.0104	0.0076
2nd quintile	–0.0044	0.0048	0.0026	0.0069	–0.0090	0.0066	–0.0069	0.0094	–0.0089	0.0089	–0.0009	0.0075
3rd quintile	0.0051	0.0046	0.0099	0.0065	0.0002	0.0065	0.0084	0.0078	0.0042	0.0079	0.0018	0.0081
4th quintile	0.0092	0.0047**	0.0045	0.0070	0.0113	0.0064*	0.0112	0.0075	0.0062	0.0081	0.0059	0.0089
Health												
Excellent	–0.0168	0.0047	–0.0140	0.0068**	–0.0196	0.0064***	–0.0102	0.0073	–0.0182	0.0084**	–0.0231	0.0088***
Very good	–0.0036	0.0036	0.0020	0.0052	–0.0094	0.0049*	–0.0010	0.0061	–0.0051	0.0066	–0.0044	0.0060
Fair	0.0250	0.0045***	0.0220	0.0068***	0.0253	0.0060***	0.0181	0.0091**	0.0228	0.0090**	0.0295	0.0067***
Poor	0.0600	0.0091***	0.0410	0.0145***	0.0705	0.0121***	0.0394	0.0209*	0.0821	0.0178***	0.0612	0.0125***
Pension type												
DC	–0.0146	0.0042***	–0.0124	0.0064*	–0.0134	0.0056**	–0.0146	0.0078*	–0.0036	0.0073	–0.0201	0.0068***
DB	0.0156	0.0041***	0.0270	0.0063***	0.0080	0.0054	0.0279	0.0067***	0.0089	0.0079	0.0082	0.0069
DB and DC	0.0264	0.0045***	0.0298	0.0066***	0.0261	0.0061***	0.0212	0.0076***	0.0320	0.0076***	0.0347	0.0077***
Tenure	0.0004	0.0001***	0.0009	0.0002***	–0.0003	0.0002	0.0005	0.0002**	0.0001	0.0003	0.0004	0.0002*
Marginal effect of unemployment rate at percentiles of the distribution												
Unemployment rate (p10)	–0.0037	0.0009***	–0.0039	0.0013***	–0.0035	0.0012***	–0.0043	0.0019**	–0.0074	0.0020***	–0.0017	0.0011
Unemployment rate (p25)	–0.0036	0.0008***	–0.0037	0.0012***	–0.0034	0.0011***	–0.0041	0.0017**	–0.0070	0.0018***	–0.0017	0.0010*
Unemployment rate (p50)	–0.0034	0.0008***	–0.0036	0.0011***	–0.0033	0.0010***	–0.0039	0.0016**	–0.0064	0.0015***	–0.0017	0.0010*
Unemployment rate (p75)	–0.0033	0.0007***	–0.0034	0.0010***	–0.0032	0.0009***	–0.0037	0.0014***	–0.0058	0.0013***	–0.0017	0.0010*
Unemployment rate (p90)	–0.0031	0.0006***	–0.0032	0.0008***	–0.0030	0.0008***	–0.0035	0.0012***	–0.0052	0.0009***	–0.0017	0.0009*
N	34,895		15,327		19,568		10,884		9,711		14,297	

NOTE: Observations at the person-year level for a sample from the Health and Retirement Study, original HRS cohort, from 1992-2010. Coeff., coefficiant. S.E., standard error. The sample consists of those observed for at least two waves between ages 50-69 and initially employed or self-employed and who live in an MSA or μSA and for whom self-reports or imputations exist for all the variables in the regression. The table reports marginal effects estimated from a weighted multinomial logit (using survey weights to make the sample nationally representative) with five job transitions observed based on recall data from one birthday to the next (stay in current job, involuntary exit to new job, voluntary exit to new job, involuntary exit to retirement, voluntary exit to employment). The table also reports standard errors clustered at the person level, with statistical significance denoted by *** (1% level), ** (5% level), and * (10% level). The unemployment rate is measured at the MSA or μSA level. See the text for more information about control variables.

Impact of the Local Unemployment Rate

We find that the local (MSA or μSA) unemployment rate has large and statistically significant effects on the job transitions we consider. These effects arise in the full sample, and at the disaggregated level they are somewhat similar for men and women and for workers of different skill levels; for some transitions they are stronger for semi-skilled than for skilled or unskilled workers. We tried to estimate multinomial logits on a smaller number of outcomes by investigating various combinations of the five outcomes listed in the boxed insert. However, likelihood ratio tests strongly reject the equality of coefficients across different combinations of outcomes (including involuntary plus voluntary exits to new jobs, involuntary exits to new jobs plus involuntary exits to retirement, and voluntary exits to new jobs plus voluntary exits to retirement). We also tried a greater number of outcomes by distinguishing between new jobs with full-time hours and those with part-time hours; we did not find that the unemployment rate had a significantly different effect on the likelihood of moving into a full-time or a part-time job regardless of whether the transition was voluntary or involuntary.

For the full sample (shown in column (A) of all tables), the MSA unemployment rate has significant negative effects on the likelihood of voluntary exit to either a new job (Table 3) or to retirement (Table 5) and positive effects, though small, on involuntary exits to both a new job (Table 2) and to retirement (Table 4). Interestingly, a higher local unemployment rate also raises the propensity to stay in the current job (Table 1), showing that older workers are insulated from the effects of business cycles and also choose to delay retirement, perhaps because of wealth effects or increased uncertainty about retirement resources.[18] These effects are all statistically significant at the 1 percent level, except for involuntary exits to a new job, where the effect is significant at the 10 percent level. It is not surprising that high unemployment raises involuntary exits, and it is interesting to note that the resulting job transition is just as likely to involve a new job as to involve retirement, though Chan and Stevens (2001, 2004) emphasize the retirement channel.

The magnitudes of the estimated effects of local unemployment are relatively important in size. For voluntary exit to a new job (Table 3), the marginal effect is –0.0027, so a 1-percentage-point increase in the MSA unemployment rate from the mean of 5.4 percent reduces the likelihood of this event by 0.27 percentage points, a 6.9 percent reduction for those turning age 56. Similarly, a 1-percentage-point increase in the local unemployment rate reduces the likelihood of a voluntary exit to retirement (Table 5) by 0.34 percentage points, while it raises the likelihood of staying in the same job (Table 1) by 0.49 percentage points.

We can put this further in perspective by evaluating the effects of the Great Recession based on our model estimates. We account for nonlinearities by calculating predicted transitions at two different unemployment rates representing the peak of the business cycle in 2007 and the trough of the recession in 2010. We find evidence of small nonlinearities, with a decrease in estimated marginal effects of 10 to 20 percent when the unemployment rate doubles.[19] Our estimates show that the increase in the unemployment rate from a low of 4.6 percent to a high of 9.6 percent resulted in a predicted increase in involuntary exits to new jobs and to retirement of 0.29 and 0.32 percentage points, respectively, and a predicted decrease in voluntary exits to new jobs and to retirement of 1.14 and 1.59 percentage points, respectively.

The remaining table columns show whether responses are heterogeneous for different types of workers. The effects of the local unemployment rate are similar for men (column (B)) and women (column (C)), except that men are somewhat more likely to stay in their current job in response to a higher unemployment rate (with a 1-percentage-point increase raising the likelihood of staying by 0.57 percentage points for men and by 0.44 points for women). In addition, involuntary exits for women are more likely to lead to a new job and for men are more likely to lead to retirement. These differences may arise because men's jobs are more remunerative on average, facilitating retirement, and perhaps because husbands lead wives in making joint retirement decisions.

The effects of local labor market conditions also vary in some ways by worker skill levels, as the unemployment rate may have different effects on skill-specific labor markets.[20] Stronger effects arise for semi-skilled workers by altering voluntary exits to retirement, while for other transitions the effect of the unemployment rate is similar across workers in different occupations.[21] A 1-percentage-point increase in the local unemployment rate raises the likelihood of staying in the current job by 0.48 for skilled workers, 0.86 for semi-skilled workers, and 0.31 for unskilled workers, while it reduces the likelihood of voluntary exits to retirement by 0.39, 0.63, and 0.17 percentage points, respectively.

Impact of Other Variables

When we compared the multinomial logit results with and without controlling for the MSA unemployment rate, we found very small differences in estimated effects of other variables. Thus, the effect of the unemployment rate is quite uniform across individuals who vary considerably in their socioeconomic characteristics.

Other statistically significant variables include the following. First, consider individual non-job characteristics. For the sample as a whole in column (A), men have statistically significant but relatively small differences in their job transitions, with the largest difference being a 1.6-percentage-point reduction in the likelihood of voluntary exits to retirement. Education has little effect on involuntary exits, while higher educational attainment is associated with an increased likelihood of voluntary exit to another job rather than to retirement (so educated workers voluntarily work longer in bridge jobs). Health has little association with taking a new job versus staying in the same job, but excellent health substantially reduces the likelihood of exiting to retirement (either involuntarily or voluntarily), while poor health substantially raises it, relative to staying in the same job.

Next, consider job characteristics, again for the sample as a whole in column (A). Blue-collar industries (agriculture/mining/construction [noted "mining" on the tables] and manufacturing/transport [noted "manufacturing" on the tables]) generate significantly more involuntary exits in total as well as more voluntary exits to retirement. White-collar industries (professional services/public administration [noted "professional" on the tables]) generate significantly fewer involuntary exits. Also, semi-skilled occupations (sales/clerical) are most likely to experience involuntary exits to retirement.

Previous research shows that employer-provided pensions can have substantial effects on the timing and manner of exits from career jobs. Here, we find that having any type of pension reduces the likelihood of involuntary exits, as pensioned jobs are probably more stable, while it also reduces the likelihood of voluntary exits to another job. This finding is consistent with evidence in Friedberg and Owyang (2005) that workers with any type of pension have longer tenure in jobs, with greater effects for workers with defined benefit pensions than for workers with only defined contribution pensions. Meanwhile, workers with defined benefit pensions are substantially more likely to exit voluntarily to retirement, especially when they are older than the plan's normal retirement age; conversely, workers with defined contribution plans are less likely to voluntarily retire, as in Friedberg and Webb (2005).

CONCLUSION

The ability of employees to exit the labor force at an age and in a manner of their choosing depends on their ability to find employment at older ages, which depends in turn on local labor market conditions. Thus, we investigate how local labor market conditions affect retirement transitions, a question that has until recently been overlooked in the retirement literature. To study this, we use data from the HRS, the first dataset to offer both a lengthy panel and rich local identifiers on a restricted basis. This level of detail allows us to estimate a multinomial logit model that distinguishes among the multifaceted paths that workers take to retirement.

We find that the local unemployment rate has statistically significant and important effects on retirement transitions. Interestingly, a higher MSA unemployment rate significantly reduces the likelihood of voluntary exits from a job, probably reflecting the corresponding difficulty of finding a new job at older ages and possible wealth effects of recessions. A higher unemployment rate also has significant but relatively small effects in terms of raising the likelihood of involuntary exits, generating about equal movements to new jobs and to retirement.

The magnitudes of the estimated effects of local unemployment are important in size. A 1-percentage-point increase in the MSA unemployment rate reduces the likelihood of voluntary exit to a new job by 6.9 percent. It also reduces the likelihood of voluntary exit to retirement by 8.3 percent, while it raises the likelihood of involuntary exit to retirement by 5.5 percent. Our findings that local labor markets influence retirement transitions, and especially phased retirement, are particularly interesting as we emerge from a severe recession that has eroded retirement portfolios. These findings shed light on how the recession's impact on labor markets will in turn affect the retirement patterns of the Baby Boom generation. ■

NOTES

1 The HRS geographic identifiers are available to qualified researchers under conditions that prevent identification of particular MSAs. The HRS, with its focus on older workers, does not allow a broad analysis of local labor markets and prime-age workers; most studies on prime-age workers (for example, Sullivan and von Wachter, 2009, and Couch and Placzek, 2010) have used unemployment records from a single state.

2 Such studies assume some sort of rigidity in wages, with downward rigidity preventing wages from adjusting fully in response to negative labor demand shocks and resulting in employment responses. For many reasons, it is likely that wages are even more downwardly rigid for older workers than for prime-age workers; see von Wachter (2007) for a thorough discussion.

3 As noted later, we tried to categorize local and industry-specific unemployment rates, similar to von Wachter's approach but at a more detailed industry level; we did not find a significant effect at the industry level. It is possible that von Wachter's state- and industry-specific unemployment rates serve as proxies for local conditions, to the extent that industries cluster locally within states.

4 Unlike their approach, we focus only on objective measures of work and nonwork. In any case, they used the empirical results incidentally, as their main focus was an analytical equilibrium search model.

5 Two other articles have analyzed local labor markets in Britain. Haardt (2006) and Disney, Ratcliffe, and Smith (2015) both estimated hazard models and found conflicting results. While Haardt found that the regional unemployment rate significantly reduces women's exits from and returns to the labor force, Disney, Ratcliffe, and Smith (2015) found that the local unemployment rate raises the hazard rate of full withdrawal from the labor force.

6 The use of CPS data in Coile and Levine (2007, 2011), which allows them to examine business cycles over a long period, necessitates a much simpler focus on retirement transitions. They define retirement as occurring when someone worked at least 13 weeks in the previous year and is out of the labor force in the March survey.

7 Note that voluntary job exits may be due to factors beyond the individual's control—for example, in case of illness or spousal unemployment—but such exits still reflect labor supply considerations rather than labor demand. In some specifications, we allow $K = 7$ outcomes by considering whether new jobs taken voluntarily involved part-time or full-time hours, but we do not find significantly different effects of the local unemployment rate on this choice. The HRS did not have a large enough sample of people leaving their jobs involuntarily and taking new jobs to statistically distinguish the effect of the local unemployment rate on part-time versus full-time new jobs.

8 The HRS is administered by the Institute for Social Research (ISR) at the University of Michigan with support from the National Institute on Aging and the SSA. Where possible, we use the RAND HRS data file, a cleaned version of the original. We have not incorporated new cohorts entering the HRS in 1998 or 2004.

9 Specific reasons reported by respondents for changes in an employment situation that encouraged them to leave are (i) their supervisor or a coworker encouraged their departure, (ii) their wage or hours were reduced, or (iii) they would have been laid off.

10 The U.S. Census Bureau has defined 940 CBSAs for the country. A CBSA consists of one or more counties or county equivalents that have at least one urban core area of at least 10,000 population, plus adjacent territory that has a high degree of economic and social integration with the core as measured by commuting ties (U.S. Office of Management and Budget, 2006). These CBSAs are divided into 363 MSAs with core areas of at least 50,000 and 577 smaller micropolitan statistical areas (μSAs). As of the 2000 Census, 82.6 percent of the population lived in MSAs, 10.3 percent in μSAs, and 7.1 percent in neither. We experimented with an alternative of using combined statistical areas where appropriate and obtained substantially similar results. Combined statistical areas are groups of CBSAs with substantial commuting ties.

11 We find that after inclusion of sample weights, the sample is indeed broadly nationally representative.

12 The HRS provides 13 industry and 17 occupation codes, derived from the 2000 Census industry and occupation codes. Based on previous literature, we group industry codes 1-2 as agriculture/construction/mining, 3-5 as manufacturing, 6-11 as professional services, and 12-13 as nonprofessional services. We group occupation codes 1-2 (managerial, professional) as skilled, 3-4 (sales, clerical) as semi-skilled, and all others as unskilled.

13 Employer-reported information available in the HRS is more accurate but is available for only about 70 percent of the sample. While Gustman and Steinmeier (2004) showed that individuals report pension information with substantial error, Chan and Stevens (2008) found that retirement responded more to one's beliefs about one's pension type, but also that, as people approached retirement, the accuracy of their information improved.

[14] Social Security earnings records, which can be used to compute Social Security wealth (SSW) and Social Security "peak value" (the discounted gain in SSW available if an individual waits to retire until SSW reaches its peak, as in Coile and Gruber, 2007) are reported for respondents who gave permission to be matched to Social Security records; these records are normally available to qualifying researchers on a restricted basis. However, any use that combines both restricted Social Security and restricted geographic data can only be undertaken onsite at the University of Michigan (ISR). In preliminary analysis at the ISR, we found that SSW peak value had a statistically significant effect on retirement, but including it did not alter estimated effects of the unemployment rate. Therefore, we did not travel again to the ISR and have reported final results without Social Security controls.

[15] We thus do not explain choices prior to the beginning of the survey; this is common to most of the retirement literature. As a consequence, we miss some early retirements, and our sample may be biased by including a disproportionate share of individuals who do not have a propensity to retire early.

[16] In contrast to our annual approach, Gustman and Steinmeier (2001/2002) tracked individuals by wave (over two years), which reduces the precision in predicting retirement since many important milestones, such as attaining age 62 or 65, or one's normal retirement age, occur on the individual's birthday.

[17] We use sample weights so the results are nationally representative.

[18] With our data it is not possible to compare this estimated effect for older workers with workers younger than age 50.

[19] For the sake of brevity, we have not reported the marginal effects estimated at other points of the distribution of the unemployment rate besides the mean. As an example of the slight nonlinearity, the marginal effect of the unemployment rate on staying in one's job is 0.0049, compared with 0.0054 at the 25th percentile value and 0.0046 at the 75th percentile values of the unemployment rate observed in our sample.

[20] We explored this further by matching workers' two-digit occupation and industry to occupation- and industry-specific unemployment rates by year and location. We had to use more aggregated geographic information and also faced the complication of substantial changes in occupation and industry coding following the 2000 Census, and we did not obtain significant results. As we noted earlier, von Wachter (2007) conducted his analysis at the state and 1-digit-industry level and found significant differences in layoff shocks across industries; it is possible that our local unemployment rates capture important industry-level shocks.

[21] A higher unemployment rate raises the likelihood of involuntary exits to new jobs for skilled and unskilled workers but does not change this likelihood for semi-skilled workers, but these effects are not statistically significant.

REFERENCES

Black, Dan; Kolesnikova, Natalia and Taylor, Lowell J. "Why Do So Few Women Work in New York (and So Many in Minneapolis)? Labor Supply of Married Women Across U.S. Cities." *Journal of Urban Economics*, January 2014, *79*, pp. 59-71; https://doi.org/10.1016/j.jue.2013.03.003.

Black, Dan and Liang, Xiaoli. "Local Labor Market Conditions and Retirement Behavior." Working Paper No. 2005-08, Boston College Center for Retirement Research, May 2005.

Callaway, Brantly. "Job Displacement of Older Workers During the Great Recession: Tight Bounds on Distributional Treatment Effect Parameters Using Panel Data." Unpublished manuscript, Vanderbilt University, October 24, 2015.

Chan, Sewin and Stevens, Ann Huff. "Job Loss and Employment Patterns of Older Workers." *Journal of Labor Economics*, April 2001, *19*(2), pp. 484-521; https://doi.org/10.1086/319568.

Chan, Sewin and Stevens, Ann Huff. "How Does Job Loss Affect the Timing of Retirement?" *B.E. Press Contributions to Economic Analysis and Policy*, May 2004, *3*(1).

Chan, Sewin and Stevens, Ann Huff. "What You Don't Know Can't Help You: Pension Knowledge and Retirement Decision-Making." *Review of Economics and Statistics*, May 2008, *90*(2), pp. 253-66; https://doi.org/10.1162/rest.90.2.253.

Coile, Courtney C. and Gruber, Jonathan. "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics*, May 2007, *89*(2), pp. 234-46; https://doi.org/10.1162/rest.89.2.234.

Coile, Courtney C. and Levine, Phillip B. "Labor Market Shocks and Retirement: Do Government Programs Matter?" *Journal of Public Economics*, November 2007, *91*(10), pp. 1902-19; https://doi.org/10.1016/j.jpubeco.2007.01.005.

Coile, Courtney C. and Levine, Phillip B. *Reconsidering Retirement: How Losses and Layoffs Affect Older Workers*. Washington, DC: Brookings Institution Press, 2010.

Coile, Courtney C. and Levine, Phillip B. "The Market Crash and Mass Layoffs: How the Current Economic Crisis May Affect Retirement." *B.E. Journal of Economic Analysis and Policy*, January 2011, *11*(1); https://doi.org/10.2202/1935-1682.2568.

Couch, Kenneth and Placzek, Dana. "Earnings Losses of Displaced Workers Revisited." *American Economic Review*, March 2010, *100*(1), pp. 572-89; https://doi.org/10.1257/aer.100.1.572.

Disney, Richard; Ratcliffe, Anita and Smith, Sarah. "Booms, Busts and Retirement Timing." *Economica*, July 2015, *82*(327), pp. 399-419; https://doi.org/10.1111/ecca.12133.

French, Eric and Jones, John B. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." *Econometrica*, May 2011, *9*(3), pp. 693-732.

Friedberg, Leora and Owyang, Michael T. "Explaining the Evolution of Pension Structure and Job Tenure." Federal Reserve Bank of St. Louis Working Paper No. 2002-022D, November 2005; https://research.stlouisfed.org/wp/2002/2002-022.pdf.

Friedberg, Leora and Webb, Anthony. "Retirement and the Evolution of Pension Structure." *Journal of Human Resources*, Spring 2005, *40*(2), pp. 281-308; https://doi.org/10.3368/jhr.XL.2.281.

Goda, Gopi Shah; Shoven, John B. and Slavov, Sita Nataraj. "Does Stock Market Performance Influence Retirement Intentions?" *Journal of Human Resources*, 2012, *47*(4), pp. 1055-81; https://doi.org/10.3368/jhr.47.4.1055.

Gustman, Alan and Steinmeier, Thomas. "Retirement Measures in the Health and Retirement Study." *Journal of Human Resources*, January 1995, *30*(1, Suppl.), pp. s57-s83; https://doi.org/10.2307/146278.

Gustman, Alan and Steinmeier, Thomas. "Retirement and Wealth." *Social Security Bulletin*, 2001/2002, *64*(2), pp. 66-91; https://www.ssa.gov/policy/docs/ssb/v64n2/v64n2p66.pdf.

Gustman, Alan and Steinmeier, Thomas. "What People Don't Know About Their Pensions and Social Security," in William Gale, John Shoven, and Mark Warshawsky, eds., *Private Pensions and Public Policies*. Washington, DC: Brookings Institution, 2004, pp. 57-125.

Gustman, Alan and Steinmeier, Thomas. "The Social Security Early Retirement Age in a Structural Model of Retirement and Wealth." *Journal of Public Economics*, February 2005, *89*(2-3), pp. 441-63; https://doi.org/10.1016/j.jpubeco.2004.03.007.

Haardt, David. "Transitions Out Of and Back To Employment Among Older Men and Women in the UK." ISER Working Paper No. 2006-20, Institute for Social and Economic Research, May 2006; https://www.iser.essex.ac.uk/research/publications/working-papers/iser/2006-20.

Hairault, Jean-Olivier; Langot, François and Zylberberg, André. "Equilibrium Unemployment and Retirement." *European Economic Review*, October 2015, *79*, pp. 37-58; https://doi.org/10.1016/j.euroecorev.2015.07.002.

Hutchens, Robert. "Phased Retirement: Problems and Prospects." Center for Retirement Research at Boston College *An Issue in Brief*, February 2007, *8*, pp. 66-91.

Maestas, Nicole. "Back to Work: Expectations and Realizations of Work After Retirement." *Journal of Human Resources*, Summer 2010, *45*(3), pp. 718-48; https://doi.org/10.3368/jhr.45.3.718.

Maestas, Nicole; Mullen, Kathleen J. and Powell, David. "The Effect of Local Labor Demand Conditions on the Labor Supply Outcomes of Older Americans." RAND Working Paper No. WR-1019, September 2013; https://www.rand.org/content/dam/rand/pubs/working_papers/WR1000/WR1019/RAND_WR1019.pdf.

Munnell, Alicia H.; Soto, Mauricio; Triest, Robert and Zhivan Natalia. "How Much Do State Economic and Other Characteristics Affect Retirement Behavior?" Unpublished manuscript, Boston College Center for Retirement Research, September 2008.

Owyang Michael T.; Piger, Jeremy M. and Wall, Howard J. "Discordant City Employment Cycles." *Regional Science and Urban Economics*, March 2013, *43*(2), pp. 367-84; https://doi.org/10.1016/j.regsciurbeco.2012.09.004.

Owyang Michael T.; Piger, Jeremy M.; Wall, Howard J. and Wheeler, Christopher H. "The Economic Performance of Cities: A Markov-Switching Approach." *Journal of Urban Economics*, November 2008, *64*(3), pp. 538-50; https://doi.org/10.1016/j.jue.2008.05.006.

Ozturk, Gulgun Bayaz and Gallo, William T. "Effect of Job Loss on Wealth Accumulation of Older Workers." CUNY School of Public Health at Hunter College, January 2013.

Ruhm, Christopher. "Bridge Jobs and Partial Retirement." *Journal of Labor Economics*, October 1990, *8*(4), pp. 482-501; https://doi.org/10.1086/298231.

Rust, John and Phelan, Christopher. "How Social Security and Medicare Affect Retirement Behavior in a World with Incomplete Markets." *Econometrica*, July 1997, *65*(4), pp. 781-831; https://doi.org/10.2307/2171940.

Sullivan, Daniel and von Wachter, Till. "Job Displacement and Mortality: An Analysis Using Administrative Data." *Quarterly Journal of Economics*, August 2009, *124*(3), pp. 1265-306; https://doi.org/10.1162/qjec.2009.124.3.1265.

Topel, Robert H. "Local Labor Markets." *Journal of Political Economy*, June 1986, *94*(3, Part 2), pp. S111-S143; https://doi.org/10.1086/261401.

U.S. General Accounting Office. *Older Workers: Demographic Trends Pose Challenges for Employers and Workers*. General Accounting Office Report GAO 02-85. Washington, DC: U.S. General Accounting Office, November 2001; http://www.gao.gov/new.items/d0285.pdf.

U.S. Office of Management and Budget. *Updates to Statistical Areas*. OMB Bulletin No. 07-01, December 18, 2006.

von Wachter, Till. "The Effect of Economic Conditions on the Employment of Workers Nearing Retirement Age." Working Paper No. 2007-25, Boston College Center for Retirement Research, December 2007.

Wall, Howard, J. "The Employment Cycles of Neighboring Cities." *Regional Science and Urban Economics*, January 2013, *43*(1), pp. 177-85; https://doi.org/10.1016/j.regsciurbeco.2012.06.008.

Model Averaging and Persistent Disagreement

In-Koo Cho and Kenneth Kasa

The authors consider the following scenario: Two agents construct models of an endogenous price process. One agent thinks the data are stationary, the other thinks the data are nonstationary. A policymaker combines forecasts from the two models using a recursive Bayesian model averaging procedure. The actual (but unknown) price process depends on the policymaker's forecasts. The authors find that if the policymaker has complete faith in the stationary model, then beliefs and outcomes converge to the stationary rational expectations equilibrium. However, even a grain of doubt about stationarity will cause beliefs to settle on the nonstationary model, where prices experience large self-confirming deviations away from the stationary equilibrium. The authors show that it would take centuries of data before agents were able to detect their model misspecifications. (JEL C63, D84)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2017, *99*(3), pp. 279-94.
https://doi.org/10.20955/r.2017.279-294

> Parameter uncertainty, while difficult, is at least a relatively well defined problem. Selecting the right model from among a variety of non-nested alternatives is another matter entirely. While there is some formal literature on this problem, I think it is safe to say that central bankers neither know nor care much about this literature. I leave it as an open question whether they are missing much. My approach to this problem while on the Federal Reserve Board was relatively simple: Use a wide variety of models and don't ever trust any one of them too much…My usual procedure was to simulate a policy on as many of these models as possible, throw out the outliers, and average the rest to get a point estimate of a dynamic multiplier path.
> **—Blinder (1998, p. 12)**

1 INTRODUCTION

The Federal Reserve is in the forecasting business. Monetary policy affects the economy with a lag, so policy must be based on forecasts. Unlike the weather, where the basic equations of fluid dynamics are well-known, no one knows what the economy's underlying equations

In-Koo Cho is a professor of economics at the University of Illinois, Urbana-Champaign, and Hanyang University and a research fellow at the Federal Reserve Bank of St. Louis. Kenneth Kasa is a professor of economics at Simon Fraser University. The authors thank two anonymous reviewers from the Federal Reserve Bank of St. Louis for several useful comments and suggestions and the National Science Foundation for financial support. The authors are grateful for the hospitality provided by the Federal Reserve Bank of St. Louis.

are. In fact, there is widespread disagreement about what the best approximation of these equations might be. By itself, disagreement is not a problem. Even weather forecasters face choices about the best approximation of the highly nonlinear equations of fluid dynamics. A commonly employed strategy for dealing with competing forecasts, both in economics and weather forecasting, is to *average* them, as the above quote from Blinder (1998) illustrates. The case for averaging rests on both solid decision-theoretic foundations and a wealth of practical experience.[1]

In Cho and Kasa (2016) we argue that model averaging confronts dangers that are not present in weather forecasting. When a weather forecaster calls for rain, his forecast does not influence the likelihood of rain. In contrast, the forecasts of macroeconomic policymakers *do* influence the likelihood of future outcomes, since policy is based on forecasts and future outcomes depend on current policy. It is precisely this feedback, or endogeneity, that creates dangers for model averaging.

Although feedback might seem to pose insurmountable hurdles to those schooled in the natural sciences, economists have in fact devised sophisticated methods for dealing with it. In particular, Lucas (1972) showed that feedback produces a fixed-point problem in the mapping between beliefs and outcomes. Solutions of this fixed-point problem are called a rational expectations equilibrium, and models that violate this fixed-point condition are said to be vulnerable to the "Lucas critique."

Unfortunately, the solution proposed by Lucas (1972) is not that useful to macroeconomic policymakers, who must decide how to weigh competing forecasts. In a rational expectations equilibrium, model averaging becomes a moot issue, since everyone is presumed to have common knowledge of the correct underlying model. As a result, there is no disagreement to average out.

So in place of Lucas (1972), Cho and Kasa (2016) pursue an approach advocated by Hansen (2014). Hansen suggests that economists populate their models with agents who behave just like econometricians who construct simple and useful provisional models and revise them as needed in response to statistical evidence. In a sense, this is not a new idea. Adaptive learning has long been used as a defense of the rational expectations hypothesis, and there are results that support this defense.[2] However, these results do not confront the issues of model uncertainty and competing forecasts. They assume agents have somehow settled on a *single* model and must merely learn its coefficients. Hansen (2014) instead emphasizes that the real problem confronting econometricians is not how to estimate a given model, but rather how to test and discriminate among competing models, all of which are recognized to be imperfect approximations of the underlying data-generating process.

As noted in the opening quote by Blinder (1998), this is a substantially more difficult problem. In fact, without a priori knowledge of the model, the results of Cho and Kasa (2016) cast doubt on the ability of agents to learn their way to a stationary rational expectations equilibrium. In particular, we show that the mere presence of a nonstationary alternative model can produce a sort of "Gresham's law" phenomenon, in which a stationary model is driven out of consideration simply because use of the nonstationary model creates conditions that support its continued use. This volatility trap could be avoided if the policymaker had dogmatic priors that ruled out the nonstationary alternative. It could also be avoided in a Lucas world, where policymakers fully understand (somehow) the endogeneity they confront. Unfortunately, neither alternative is persuasive. Instead, using the results of Cho and Kasa (2015), we argue that in the presence of potentially misspecified endogeneity it might be preferable to avoid model averaging altogether and instead base forecasts on a process of specification testing and *model selection*.

The results in Cho and Kasa (2016) were based on the previous work of Evans et al. (2013). Following Evans et al. (2013), we assumed the competing models are in the mind of a single agent. However, if this is the case, one wonders why the agent does not just expand the model space to include a single model that nests both. As noted by Timmermann (2006), nesting does not occur in practice, partly due to degrees of freedom considerations but mainly because competing models are based on differing information sets and heterogeneous prior beliefs among forecasters.

In the current paper, we extend the results of Cho and Kasa (2016) to this more relevant decentralized environment, where competing models reflect the beliefs of distinct agents. This multiple-agent setting raises some subtle issues. Our previous analysis focused on the policymaker's problem. He had to decide how to weigh the two forecasts. Here, with multiple agents, we must also consider the problem of the forecasters. Do they know their forecasts are being used by the policymaker, and are they aware the data might be endogenous? Do they know they are competing with another forecaster? If so, what are their beliefs about their rival's model? Would forecasters ever have an incentive to revise their beliefs about either their own model or another rival's model?

We show that Gresham's law of model averaging continues to apply in this more realistic decentralized setting. The agent using the nonstationary/drifting parameters model eventually dominates. Moreover, both forecasters maintain their beliefs about both their own model and their rival's model. This is interesting because *both* agents have misspecified models. However, the nature of their misspecifications turns out to be very difficult to detect statistically in finite samples.

In particular, the agent with the drifting parameters model has misspecified beliefs about his own model. Since the underlying model features constant parameters, the actual data-generating process is stationary. However, when expectational feedback is strong, mean reversion is weak, and so the agent with the drifting parameters model confronts the problem of testing the null of a unit root against a local alternative. These tests are well-known to have low power. In fact, we show that it would take *centuries* of data before the agent could reject his model with any confidence.

The agent with the constant parameters model has a correctly specified model. However, his beliefs about his *rival's* model are misspecified. In particular, he rationalizes the ongoing parameter variation in his rival's model as reflecting noise around a constant parameter value. In reality, his rival's model features mean-reverting dynamics that pull the estimates toward the rational expectations equilibrium. However, as before, this mean reversion is quite weak, and so the agent with the constant parameters model confronts the problem of testing the null of i.i.d. fluctuations against the alternative of weak autocorrelation. Again, we find that it would take centuries of data before the agent could reject his beliefs about his rival's model with any confidence.

Although we frame our discussion within the context of a specific model, our results are of broader significance. Economists beyond a certain age will likely recall an era that is sometimes called the "unit root wars." During the 1980s and 1990s there were heated debates among economists about whether macroeconomic time series were "trend stationary" or "difference stationary." Econometricians designed ever more esoteric procedures to test these hypotheses. Although echoes of these debates still resound, for the most part the unit root wars have died out, without a clear victor. At the end of the day, what the unit root wars taught us is that it is very difficult to say anything with much confidence about stationarity. We simply do not have enough data.[3]

Although lack of a clear victor may have been disappointing to the participants, applied economists have largely taken the advice of Cochrane (1991) by avoiding the question altogether. Cochrane argues that the practically relevant question is which asymptotic distribution provides the better approximation to the actual (finite-sample) distribution, which is a question that must be decided on a case-by-case basis.

Our results sound a note of caution about this pragmatic approach to the question of stationarity. Cochrane's advice is based on the perspective of what Hansen (2014) calls the "outside econometrician," meaning someone whose actions do not influence the system he is trying to learn about. The agents in our model are engaged in their own unit root war, and perhaps not surprisingly, they find that they cannot settle it either with the available data. However, they are *not* outside econometricians. Their beliefs about stationarity influence the data-generating process. We show that doubts about stationarity can produce outcomes that statistically confirm those doubts. This is a far more pessimistic conclusion than Cochrane's. In a self-referential world, persistent belief differences about stationarity actually matter.

The remainder of the paper is organized as follows. The next section outlines a simple model that conveys the basic idea. The essential feature of this model is that current outcomes depend on beliefs about future outcomes. We assume there are two forecasters who share a common reduced-form model but have differing beliefs about parameter stability. One thinks the parameters are stable; the other thinks parameters drift. We show how the forecasters can revise their beliefs using the Kalman filter. Section 3 considers the problem of the policymaker who must decide how to weigh the competing forecasts. Like Blinder in the opening quote, we assume he hedges his bets by averaging them, with weights that are recursively updated based on historical relative performance. We state our Gresham's law result, which shows that eventually the time-varying parameters (TVP) forecaster must dominate. We do not formally prove the result but merely provide the intuition. Section 4 turns to the problem of the forecasters and asks whether they would have any reason to modify or reject their models. Since they both have misspecified models, with an infinite sample they will both be able reject their models. However, the empirically relevant question is how long it will take. We present Monte Carlo evidence that suggests it would take centuries. In the meantime, prices are substantially more volatile than they would otherwise be. Section 5 provides a brief discussion of our results and attempts to place them in context. Finally, Section 6 offers a few concluding remarks.

2 A SIMPLE MODEL

We begin by outlining a special case of the model considered by Cho and Kasa (2016), in which a price at time t, p_t, is determined according to

$$p_t = \delta + \alpha E_t p_{t+1} + \sigma \epsilon_t, \quad (1)$$

where $\alpha \in (0,1)$ is a discount rate. This is a key parameter in the ensuing analysis. It determines the strength of expectational feedback. For simplicity, the ϵ_t shock is assumed to be Gaussian white noise. This model has been a workhorse in the macroeconomic learning literature. It can be interpreted as an asset-pricing model with constant fundamentals. The unique stationary rational expectations equilibrium is

$$p_t = \frac{\delta}{1-\alpha} + \sigma \epsilon_t. \quad (2)$$

Hence, rational expectations predicts that prices exhibit i.i.d. fluctuations around a fixed mean with a variance of σ^2.

2.1 Learning

Note that rational expectations is an *equilibrium* concept. It says nothing about how, or even whether, this equilibrium could ever be attained. The original architects of the rational expectations revolution (John Muth, Robert Lucas, Edward Prescott, and Thomas Sargent) defended it by arguing that a process of adaptive, out-of-equilibrium learning would eventually produce convergence to a rational expectations equilibrium. Applications assume this process of learning has already taken place. However, formal analysis of this conjecture did not begin until the 1980s. Due to the presence of feedback, which makes the data endogenous, addressing the convergence question is not a simple matter of applying the law of large numbers.

There is now a well-developed literature that provides conditions under which a wide variety of economic models can be expected to converge to rational expectations equilibria.[4] For the particular model in equation (1), the crucial restriction is that $\alpha < 1$. That is, feedback cannot be too strong. Fortunately, this imposes no additional restrictions in this case, since theory requires $\alpha < 1$.

Our paper questions these existing convergence results. These results are based on imputing dogmatic priors to agents. The typical case presumes agents are convinced that parameters are constant; they are just not sure what the constants are. They rule out a priori the possibility that parameters drift. A more recent literature, based on so-called "constant gain" learning algorithms, presumes agents have a dogmatic prior that parameters drift. Here agents are not even given the chance to learn the constant parameter rational expectations equilibrium, because their priors say this is a zero probability event.

Here, agents are less dogmatic. In particular, the policymaker is open-minded. He puts positive weight on both possibilities and then revises his beliefs using Bayes' rule as the data come in. One might suspect, based on a naive application of "grain of truth" sort of arguments, that eventually he will learn the constant parameter rational expectations equilibrium since his prior puts positive weight on this possibility. This will indeed be the case if feedback is relatively weak. However, we show that if $\alpha > 1/2$, which is the empirically relevant case, the policymaker's beliefs will eventually converge to the time-varying parameters model. The usual grain-of-truth argument does not apply here because *both* models are misspecified due to the presence of feedback.

To be more precise, suppose there are two forecasters who share the following state-space model for prices:

$$p_t = \beta_t + \sigma\epsilon_t \tag{3}$$

$$\beta_t = \beta_{t-1} + \sigma_v v_t, \tag{4}$$

where it is assumed that $\text{cov}(\epsilon,v) = 0$. Note that the rational expectations equilibrium is a special case of this, with

$$\sigma_v = 0 \text{ and } \beta = \frac{\delta}{1-\alpha}.$$

For now, suppose the beliefs of the forecasters about parameter stability are dogmatic. Later, in Section 4, we allow the forecasters to question their priors. In particular, suppose one forecaster, whom we call $\mathcal{M}_0$, believes parameters are constant, that is,

$$\mathcal{M}_0\text{: } \sigma_v^2 = 0,$$

while the other forecaster, $\mathcal{M}_1$, is convinced parameters are time-varying, that is,

$$\mathcal{M}_1\text{: } \sigma_v^2 > 0.$$

Both forecasters use the Kalman filter to revise their beliefs about β_t:

$$\beta_{t+1}(i) = \beta_t(i) + \left(\frac{\Sigma_t(i)}{\sigma^2 + \Sigma_t(i)}\right)\left(p_t - \beta_t(i)\right). \tag{5}$$

The only difference is in how they update their beliefs about the variance of β_t:

$$\Sigma_{t+1}(0) = \Sigma_t(0) - \frac{\Sigma_t^2(0)}{\sigma^2 + \Sigma_t(0)} \tag{6}$$

$$\Sigma_{t+1}(1) = \Sigma_t(1) - \frac{\Sigma_t^2(1)}{\sigma^2 + \Sigma_t(1)} + \sigma_v^2. \tag{7}$$

Note that the parameter update, equation (5), assumes that β_t is based on time-$(t-1)$ information. This assumption is made to avoid simultaneity between beliefs and observations.[5]

It turns out that both forecasters' estimates converge to the same (rational expectations) value, $\beta_t(i) \to \delta/(1-\alpha)$. This is not too surprising since they are both forecasting the same thing using the same data and the same reduced-form model. However, there is an important difference in the nature of the convergence. The estimates of $\mathcal{M}_0$ converge in a relatively strong sense, based on the law of large numbers. As the sample size grows, it becomes increasingly unlikely that $\beta_t(0)$ will deviate from $\delta/(1-\alpha)$ by any given amount. In contrast, the estimates of $\mathcal{M}_1$ converge in a weaker, distributional sense, based on the central limit theorem. His estimates never settle down. They fluctuate persistently around the rational expectations value. These persistent fluctuations reflect this forecaster's desire to remain alert to potential parameter instability.

One might suspect that since the underlying structural parameters (δ,α) are constant, $\mathcal{M}_0$ is "right" and $\mathcal{M}_1$ is "wrong," so that if someone were able to cast and revise a (weighted) vote on who's right, eventually $\mathcal{M}_0$ would win. In fact, precisely the opposite occurs.

3 MODEL AVERAGING

We now turn to the problem of the policymaker. The policymaker does not construct models himself. He is presented with competing forecasts and must decide how to use them. Like Alan Blinder in the opening quote, we assume he does this by averaging them. Specifically, let π_t denote the current probability assigned by the policymaker to $\mathcal{M}_1$, the TVP model, and let $\beta_t(i)$ denote the current parameter estimate for $\mathcal{M}_i$. The policymaker's time-t forecast becomes

$$\mathrm{E}_t p_{t+1} = \pi_t \beta_t(1) + (1-\pi_t)\beta_t(0).$$

Substituting this into the actual law of motion for prices in equation (1) implies that the forecasters' parameter estimates evolve according to

$$\beta_{t+1}(i) = \beta_t(i) + \left(\frac{\Sigma_t(i)}{\sigma^2 + \Sigma_t(i)}\right)\left\{\left[\delta + \alpha\left[\pi_t\beta_t(1) + (1-\pi_t)\beta_t(0)\right] - \beta_t(i)\right] + \sigma\epsilon_t\right\}. \tag{8}$$

Notice the presence of feedback here. The evolving beliefs of the forecasters depend on the beliefs of the policymaker. The more confidence the policymaker has in $\mathcal{M}_i$, the more likely $\mathcal{M}_i$ will prevail. For now, we suppose the forecasters are unaware of this feedback when constructing their models.

Since $\beta_t(i) \to \delta/(1-\alpha)$ for $\mathcal{M}_i$, where $i = 1, 2$, the only real question is what happens to π_t. We assume the policymaker is a Bayesian. He views the current value of π_t as a prior and then updates it using Bayes' rule,

$$\frac{1}{\pi_{t+1}} - 1 = \frac{A_{t+1}(0)}{A_{t+1}(1)}\left(\frac{1}{\pi_t} - 1\right), \tag{9}$$

where

$$A_t(i) = \frac{1}{\sqrt{2\pi\left(\sigma^2 + \Sigma_t(i)\right)}}\exp\left[-\frac{\left(p_t - \beta_t(i)\right)^2}{2\left(\sigma^2 + \Sigma_t(i)\right)}\right] \tag{10}$$

is the time-t likelihood of model i.

The ultimate fate of this economy has now been converted to a mathematical problem. Equations (6), (7), (8), (9), and (10) are a system of seven nonlinear stochastic difference equations in $(\pi_t, \beta_t(i), \Sigma_t(i), A_t(i))$, where $i = 1, 2$. Solving even one nonlinear stochastic difference equation can be challenging, so this would seem to be a hopeless endeavor. The key to making this system tractable is to exploit the fact that subsets of the variables evolve on different timescales. By appropriately averaging over each subset, we can simplify the analysis to one of studying the interactions between lower dimensional subsystems. In particular, variables that operate on a relatively slow timescale can be fixed at their current values, while variables that operate on a relatively fast timescale can be fixed at the means of their stationary distributions. In Cho and Kasa (2016) we show that this economy features a hierarchy of *four* timescales. The data operate on a relatively fast calendar timescale. Estimates of the TVP model evolve on a slower timescale determined by the parameter innovation variance. Estimates of the constant parameter model evolve even more slowly, on a timescale determined by the inverse of the historical sample size. Finally, the model weight, π_t, evolves on a variable timescale but spends most of its time in the neighborhood of either 0 or 1, where it evolves on a timescale that is even slower than that of the constant parameter model.

Exploiting these timescale separation methods, Cho and Kasa (2016) prove the following result.

Theorem 3.1. $\forall \epsilon > 0, \forall T \geq 1$, *define*

$$T_1^{\epsilon} = \#\left\{ t \leq T \middle| \pi_t \geq 1 - \epsilon \right\}$$

as the number of periods $t \leq T$ when the policymaker believes in $\mathcal{M}_1$ with probability more than $1 - \epsilon$. If $\pi_0 \neq 0$ *and* $1/2 < \alpha < 1$, *then* $\forall \epsilon > 0$,

$$\lim_{\sigma_v^2 \to 0} \lim_{T \to \infty} \mathrm{E} \frac{T_1^{\epsilon}}{T} = 1.$$

This result says that as long as feedback is sufficiently strong ($\alpha > 1/2$) and assumed parameter drift is sufficiently weak ($\sigma_v^2 \to 0$), the TVP will eventually dominate. Note the order of limits is important here. By letting $T \to \infty$ first, what we are doing is comparing asymptotic distributions as the parameter σ_v^2 changes. The same strategy is used when using stochastic stability arguments to select among multiple Nash equilibria (Kandori, Mailath, and Rob, 1993).[6] Of course, from a practical standpoint, the relevant question is what happens for strictly positive values of σ_v^2. Using simulations, Cho and Kasa (2016) show that convergence is continuous and that values of σ_v^2 that generate significant steady-state "excess volatility" are associated with a value of π_t well in excess of 1/2. The following section provides additional results along these lines.

The intuition for why the TVP model eventually dominates is the following: When the weight on the TVP model is close to 1, the world is relatively volatile (due to feedback). This makes the constant parameters model perform relatively poorly since it is unable to track the feedback-induced time-variation in the data. Of course, the tables are turned when the weight on the TVP model is close to zero. Now the world is relatively tranquil, and the TVP model's additional noise puts it at a disadvantage. However, as long as this noise isn't too large, the TVP model can exploit its ability to respond to rare sequences of shocks that generate "large deviations" in the estimates of the constant parameters model. That is, during tranquil times, the TVP model is lying in wait, ready to pounce on large-deviation events. These events provide a foothold for the TVP model, which due to feedback allows it to regain its dominance.

Interestingly, the TVP model asymptotically dominates here because it is better able to react to the volatility that it itself creates. Although $\mathcal{M}_1$ is misspecified from the perspective of the rational expectations equilibrium, this equilibrium must be *learned* via some adaptive process. Our result shows that this learning process can be subverted by the mere presence of a misspecified alternative, even when the correctly specified model would converge if considered in isolation. This result therefore echoes the conclusions of Sargent (1993), who notes that adaptive learning models often need a lot of "prompting" before they converge. Elimination of misspecified alternatives can be interpreted as a form of prompting.

4 PERSISTENT DISAGREEMENT

The policymaker in the previous analysis is quite sophisticated. Although he is ignorant of the underlying structural model, he is aware of his own ignorance and recognizes that in the presence of model uncertainty it is wise to give all models a chance. Moreover, he is not dogmatic. As evidence accumulates, he revises his beliefs.

The same cannot be said of the forecasters. Although they revise coefficient estimates, they dogmatically adhere to their beliefs about parameter stability. Moreover, they operate in a vacuum, in the sense that they are totally unaware of each other's presence and ignore the fact that the policymaker is actually using their forecasts. Here we extend the analysis of Cho and Kasa (2016) by assuming the forecasters are a bit more sophisticated.[7]

In doing this, we assume the forecasters follow Neyman-Pearson principles. That is, they stick with a null model unless sufficient evidence mounts against it. This is not as schizophrenic as it might seem. Bayesian decision theory presumes agents have full confidence in their priors. If they didn't, they would have different priors. Policymakers must decide how best to use the information they receive. In doing this, it makes sense to adhere to Bayesian principles. However, agents engaged in the construction of economic models confront a less structured and more open-ended problem. They must remain alert to the possibility that *all* current models are misspecified and be prepared to expand their priors in response. In this environment, we think Neyman-Pearson behavior makes sense.[8]

As before, we can summarize the beliefs of $\mathcal{M}_0$ and $\mathcal{M}_1$ as a pair of perceived state space models. The main difference is that now we must include each forecaster's belief about his rival's model. These beliefs take the form of conjectures about the other forecaster's beliefs about stationarity. $\mathcal{M}_1$ thinks $\mathcal{M}_0$ uses a constant parameter model, while $\mathcal{M}_0$ thinks $\mathcal{M}_1$ uses a *random* (not drifting) coefficients model with i.i.d. fluctuations around a constant mean. Specifically, the perceived observation equation for $\mathcal{M}_0$ is

$$p_t = (1-\pi_t)\beta_t(0) + \pi_t\overline{\beta}_t(1) + \sigma\epsilon_t, \tag{11}$$

while the state transition equation is

$$\begin{array}{rcl} \beta_t(0) & = & \beta_{t-1}(0) \\ \overline{\beta}_t(1) & = & \overline{\beta}_{t-1}(1) \end{array},$$

where an overbar is used to represent an agent's belief about the other agent's model. To simplify notation, define

$$y_t(0) = \begin{bmatrix} p_t \\ \beta_t(1) \end{bmatrix}, \quad \xi_{t-1}(0) = \begin{bmatrix} \beta_t(0) \\ \overline{\beta}_t(1) \end{bmatrix}, \quad H_t = \begin{bmatrix} (1-\pi_t) \\ \pi_t \end{bmatrix}, \quad Q(0) = 0$$

and

$$R(0) = \mathrm{cov}\begin{bmatrix} \sigma\epsilon_t \\ \epsilon_{1,t} \end{bmatrix}.$$

Similarly, we can write the perceived observation equation of $\mathcal{M}_1$ as

$$p_t = (1-\pi_t)\overline{\beta}_t(0) + \pi_t\beta_t(1) + \sigma\epsilon_t \tag{12}$$

and his perceived state transition equation as

$$\overline{\beta}_t(0) = \overline{\beta}_{t-1}(0) \tag{13}$$

$$\beta_t(1) = \beta_{t-1}(1) + \epsilon_{v,t}. \tag{14}$$

Define

$$y_t(1) = \begin{bmatrix} p_t \\ \beta_t(0) \end{bmatrix}, \xi_{t-1}(1) = \begin{bmatrix} \overline{\beta}_t(0) \\ \beta_t(1) \end{bmatrix}, \text{and } Q(1) = \mathrm{cov}\begin{bmatrix} 0 \\ \epsilon_{v,t} \end{bmatrix}$$

and

$$R(1) = \mathrm{cov}\begin{bmatrix} \sigma\epsilon_t \\ \epsilon_{0,t} \end{bmatrix}.$$

As before, the evolving conditional mean and variance of the posteriors are described by the Kalman filter:

$$\xi_t(i) = \xi_{t-1} + P_{t-1}(i)H_t\left(H_t'P_{t-1}H_t + R(i)\right)^{-1}\left(y_t - H_t'\xi_{t-1}\right) \tag{15}$$

$$P_t(i) = P_{t-1}(i) - P_{t-1}(i)H_t\left(H_t'P_{t-1}(i)H_t + R(i)\right)^{-1}H_t'P_{t-1}(i) + Q(i). \tag{16}$$

Note that the forecasters are assumed to know the policymaker's current model weight, π_t.

The Kalman filter generates a sequence of hidden state estimates that can then be substituted into the perceived observation equations to generate sequences of price forecasts:

$$\begin{aligned}\hat{p}_t(0) &= (1-\pi_t)\beta(0)_t + \pi_t\overline{\beta}(1)_t \\ \hat{p}_t(1) &= (1-\pi_t)\overline{\beta}(0)_t + \pi_t\beta(1)_t.\end{aligned} \tag{17}$$

As before, the policymaker takes these two forecasts and averages them:

$$\hat{p}_t = (1-\pi_t)\hat{p}_t(0) + \pi_t\hat{p}_t(1). \tag{18}$$

The actual time-t price is then determined by the (unknown) structural model in equation (1):

$$p_t = \delta + \alpha\hat{p}_t + \sigma\epsilon_t. \tag{19}$$

Notice that if equations (17) are substituted into (18), which is then substituted into the actual price equation (19), we observe that the agents' models in (11) and (12) suffer from an additional form of misspecification, since each forecaster fails to recognize that his own forecast embodies a form of model averaging in its response to the other forecaster's forecast. However, this misspecification disappears in the limit.

It turns out that if $\mathcal{M}_0$ and $\mathcal{M}_1$ do not reject their models, our previous Gresham's law result continues to apply: $\pi_t \to 1$, and the economy is plagued by self-confirming volatility.[9] However, because the two models are now more symmetric, the rate of convergence is slower. Hence, the key question is whether the two forecasters would have any reason to reject their models. Since both models are misspecified, we know that with an *infinite* sample they will eventually reject them. However, we also know the sun will vaporize our planet before this occurs, so the real question is how long it will take before they reject them.

To address this question we perform a simple Monte Carlo experiment. After a fixed interval, we allow each forecaster to test the specification of his model.[10] In the case of $\mathcal{M}_0$, this involves testing the null hypothesis that the error term in his model is i.i.d. against the alternative of (first-order) autocorrelation. The persistent mean-reverting learning dynamics of $\mathcal{M}_1$ make the alternative true. In the case of $\mathcal{M}_1$, this involves testing the null hypothesis that $\beta_t(1)$ is a random walk, against the alternative of stationarity. His own learning dynamics, along with the stationarity of the underlying model, again make the alternative true.[11]

Table 1 reports the results. We consider two sample lengths, $T = 100$ and $T = 600$. We set $\alpha = 0.95$ and $\sigma^2 = 0.10$, which suggests an annual time unit. The important parameter is σ_v^2, which reflects $\mathcal{M}_1$'s beliefs about parameter drift. Our theorems pertain to the limit, as $\sigma_v^2 \to 0$, but we know that at this limit the two models are equivalent and no extra volatility arises. However, if σ_v^2 is too big, it becomes unlikely that convergence to $\pi = 1$ occurs. The only way to proceed is to try out values and see what happens. Table 1 assumes $\sigma_v^2 = 0.0005$. Note, this is orders of magnitude smaller than σ^2. Good practice recommends reducing significance levels as the sample size increases. However, in Table 1 we assume the forecasters maintain a 5 percent significance level, even when the sample size increases to $T = 600$. This biases the results toward rejection. Finally, the numbers in Table 1 report averages across $N = 600$ replications. Each replication is initialized randomly in the neighborhood of $\pi \approx 1/2$.

Table 1

Monte Carlo Simulations

T	Reject(0)	Reject(1)	PR($\pi = 1$)	var(P)
100	0.022	0.492	0.358	0.109
600	0.480	0.221	0.705	0.153

NOTE: Results are averages across $N = 600$ replications. Parameters: $\sigma^2 = 0.10$; $\sigma_v^2 = 0.0005$; $\alpha = 0.95$.

From the second column, we see that when $T = 100$, $\mathcal{M}_0$ has no incentive to reject his model. In fact, the rejection probability is less than the size of the test! (which likely reflects sampling variability). On the other hand, $\mathcal{M}_1$ has nearly a 50 percent chance of rejecting his model. This discrepancy is explained by the fourth column, which reports the proportion of time that $\pi_t \in (0.95, 1)$ at the end of the sample. Evidently, there is little tendency for $\pi_t \rightarrow 1$ when sample sizes are relatively small. This is not surprising. When $T \approx 100$, the implicit gain of $\mathcal{M}_0$'s updated equation is not that different from $\mathcal{M}_1$'s. When $\pi_t \approx 0$, it becomes relatively easy for $\mathcal{M}_1$ to reject his model, since the data are being primarily generated by $\mathcal{M}_0$'s model, which features constant parameters. Still, a 50 percent rejection probability is nothing to write home about since in fact his model is misspecified. Finally, the last column reports the variance of prices. Remember, in a rational expectations equilibrium the variance would be 0.10. Although the resulting variance is significantly larger than assumed parameter drift, it is only about 10 percent larger than the rational expectations value. Again, this reflects the fact that in small samples there is a high probability of being in the $\pi = 0$ state.

Things become more interesting as T increases. From the second row we see that when $T = 600$, $\mathcal{M}_1$'s rejection probability actually *decreases* to 0.221. This happens because the likelihood that $\pi_T \approx 1$ almost doubles—to 0.705. At the same time, $\mathcal{M}_0$ now finds it easier to reject his model. Again, this is because the other guy's model is more likely to be generating the data. Observe that prices are now considerably more volatile. The variance exceeds the rational expectations value by more than 50 percent.

The results in Table 1 suggest that with reasonably high probability, disagreements about important aspects of models can persist for hundreds of years. Granted, $600 \neq \infty$, so an economic theorist might not find these results convincing. However, to policymakers and market participants, 600 years is an eternity.

5 DISCUSSION

Here we briefly attempt to place our results in context and respond to a potential criticism of the way we have modeled beliefs about parameter instability.

5.1 Merging

Economists have long been interested in the question of whether rational individuals can "agree to disagree." The classic reference is Aumann (1976). Aumann proved that disagreement is impossible, even when agents have different information, as long as they share a common prior and their posteriors (and their prior) are common knowledge. Blackwell and Dubins (1962) show that common priors are to be expected as long as agents initially agree on what is possible and what is impossible (i.e., their priors are mutually absolutely continuous). Kalai and Lehrer (1993) used these results to show that agents can learn to play Nash equilibria.

Our results are more closely aligned with recent work that introduces "frictions" in an effort to overcome these classic asymptotic merging results. Fudenberg and Levine (1993) show that disagreement can persist about off-equilibrium path events. Kurz (1994) shows that disagreement can persist when agents consider only a limited set of moments. Acemoglu, Chernozhukov, and Yildiz (2016) show that disagreement about underlying state variables can persist when agents' models are not identified. Perhaps most closely related to our work, Esponda and Pouzo (2016) show that disagreement can persist when agents' models are misspecified.

The common thread running through all this prior work is the presumption that agents have access to infinite samples. Persistence means *forever*. Clearly, it is of interest to know when disagreement will never disappear. However, that doesn't mean disagreements that take centuries to resolve are uninteresting or unimportant. We think finite samples are an interesting friction. Perhaps the reason it hasn't been more thoroughly studied is simply that results will likely be less precise, being necessarily probabilistic and more naturally stated using Neyman-Pearson language. Nonetheless, there are results on learning rates and large deviations that could profitably be used to characterize sets of models that agents are unable to discriminate among with a given sample size. Rather than speak of merging to a common singleton model, the task would be to characterize statistical "equivalence classes" of models. Although there is a recent literature along these lines on partial identification within econometrics, it has yet to make its way to the endogenous data setting of macroeconomics. The work of Hansen and Sargent (2008) on robustness and detection error probabilities provides the closest example of what we have in mind.

5.2 Stability

We suspect most Bayesians would argue that our policymaker was asking for trouble from the beginning, since his prior was nonconvex. It consisted of two discrete points, $\sigma_v^2 = 0$ and some other strictly positive value. Repeated switching between $\pi = 0$ and $\pi = 1$ would be interpreted as evidence that the data prefer an intermediate value. Why not let the agents estimate σ_v^2 directly? True, his filtering problem would then be nonlinear, but that's what computers are for. We have two responses to this critique. First, we agree that if a *single agent* is viewed as both formulating models and using them, then a nonconvex prior makes little sense. However, as discussed previously, actual policy environments are more decentralized and feature distinct agents undertaking distinct tasks using distinct information and their own distinct priors about how the economy works. Nesting these heterogeneous beliefs within a single convex "hypermodel" is wishful thinking. Second, the claim that allowing agents to estimate σ_v^2 directly would avoid the volatility trap ignores the fact that the data are endogenous here. In Cho and Kasa (2016), we pursue an idea from Evans and Honkapohja (1993) and compute

Nash equilibrium values of σ_v^2. That is, we look for values of σ_v^2 that are best responses to themselves. Not surprisingly, $\sigma_v^2 = 0$ is a Nash equilibrium. If the world is tranquil, then constant parameter models are optimal. However, we also show that there exist strictly positive values of σ_v^2 that are (stable) Nash equilibria, even when agents are allowed to incrementally adjust their own estimates of σ_v^2. When the world is turbulent, it is individually rational to use a time-varying parameter model, even if you know that if everyone else were to use a constant parameter model a more stable Pareto superior outcome would result.

6 CONCLUSION

One of the occupational hazards of being an economist is enduring the inevitable jokes about disagreement among economists. We've all heard the one about economists being laid end-to-end and still not reaching a conclusion, or the more recent one about economics being the only field where two people can get a Nobel prize for saying opposite things. Like most jokes, these jokes have an element of truth to them. Economists *do* disagree with each other. However, most economic models pretend that they don't.

Does this disagreement matter? Physicists and biologists disagree too, so why should we care whether economists disagree? In this paper we've argued there is an important difference between economics and the natural sciences. Disagreement among physicists presumably doesn't alter the laws of physics. Economics isn't like that. Competing economic theories, when put into practice, have the capacity to create their own self-confirming reality. This paper discussed one particular example involving beliefs about parameter stability, but there are no doubt many others.

Given that the stakes are high, what is the best strategy for coping with this disagreement? We think our paper offers two main lessons. First, we think that examples like this should tip the scales back in favor of theory in economics. In this era of big data and machine learning, one often hears the argument that economic theory has become passé. Why not let the data decide? Our paper shows why these arguments are naive. Letting the data decide is a fine strategy when the data are exogenous, but that's not the case in economics. The agents in our economy suffer bad outcomes because they do not fully understand the nature of the endogeneity they confront. Letting the data decide produces the wrong decision. If a bright young theorist came along and suggested that learning itself might generate parameter instability, the self-confirming volatility trap could be avoided. Second, we think our paper tips the scales in favor of model selection, as opposed to model averaging. Although most econometrics texts outline methods for testing and selecting models, they do so apologetically, with the warning that such procedures lack coherent decision-theoretic foundations. Bayesian decision theory tells us that we should hedge our bets by averaging across models. Our paper shows why model selection might not be such a bad idea after all. The basic problem with averaging is that it forces models to *compete*. Competition is great when the rules are fixed and fair. But here, the time-varying parameter model can effectively alter the rules of the game in its own favor. Again, if econometricians fully understood the endogeneity they confronted, this wouldn't be an issue. When priors are correctly specified, Bayesian methods produce good outcomes. When they're not, Neyman-Pearson methods begin to look better. ■

NOTES

[1] See, e.g., Bates and Granger (1969) and Timmerman (2006).

[2] See Evans and Honkapohja (2001) for a comprehensive survey.

[3] This is not an issue of the number of observations but rather of the length of the period. Perron and Shiller (1985) show that one cannot increase statistical power by sampling a given period more frequently.

[4] Evans and Honkapohja (2001) provide the definitive summary of this literature.

[5] See Evans and Honkapohja (2001) for further discussion.

[6] Taking limits in the reverse order would be uninteresting, since the models become identical as $\sigma_v \to 0$.

[7] As noted in the Introduction, if we assume the competing models are in the mind of the policymaker himself, then the previous analysis makes more sense. This is how Evans et al. (2013) interpret their analysis. However, actual forecasting and policy environments are more decentralized and feature a division of labor between policymakers and model builders.

[8] We are not alone. For example, Gilboa, Postlewaite, and Schmeidler (2008) argue that Neyman-Pearson behavior is actually more consistent with recent developments in decision theory than is Bayesian decision theory. We should also note that Bayesian practitioners are aware of the importance of testing priors against the data (e.g., Geweke, 2010). These efforts blur the distinction between Bayesian and frequentist behavior.

[9] A formal proof is available upon request.

[10] Note, we assume they do this retrospectively, not repeatedly in real time, as in the "monitoring structural change" approach of Chu, Stinchcombe, and White (1996). Hence, standard testing procedures can be applied.

[11] Note, $\mathcal{M}_0$ tests his beliefs about his rival's model, while $\mathcal{M}_1$ tests his beliefs about his own model. This imparts maximal power to the tests, since $\mathcal{M}_0$'s beliefs about his own model are correct, while $\mathcal{M}_1$'s beliefs about $\mathcal{M}_0$'s model are also correct.

REFERENCES

Acemoglu, Daron; Chernozhukov, Victor, and Muhamet, Yildiz. "Fragility of Asymptotic Agreement under Bayesian Learning." *Theoretical Economics*, January 2016, *11*(1), pp. 187-225; https://doi.org/10.3982/TE436.

Aumann, Robert J. "Agreeing to Disagree." *Annals of Statistics*, November 1976, *4*(6), pp. 1236-39; https://doi.org/10.1214/aos/1176343654.

Bates, J.M. and Granger, C.W.J. "The Combination of Forecasts." *Journal of the Operational Research Society*, December 1969, *20*(4), pp. 451-68; https://doi.org/10.1057/jors.1969.103.

Blackwell, David and Dubins, Lester. "Merging of Opinions with Increasing Information." *Annals of Mathematical Statistics*, September 1962, *33*(3), pp. 882-86; https://doi.org/10.1214/aoms/1177704456.

Blinder, Alan S. *Central Banking in Theory and Practice*. Cambridge: MIT Press, 1998.

Cho, In-Koo and Kasa, Kenneth. "Learning and Model Validation." *Review of Economic Studies*, January 2015, *82*(1), pp. 45-82; https://doi.org/10.1093/restud/rdu026.

Cho, In-Koo and Kasa, Kenneth. "Gresham's Law of Model Averaging." Working Paper No. 6, Simon Fraser University, 2016.

Chu, Chia-Shang James; Stinchcombe, Maxwell and White, Halbert. "Monitoring Structural Change." *Econometrica*, September 1996, *64*(5), pp. 1045-65; https://doi.org/10.2307/2171955.

Cochrane, John H. "A Critique of the Application of Unit Root Tests." *Journal of Economic Dynamics and Control*, April 1991, *15*(2), pp. 275-84; https://doi.org/10.1016/0165-1889(91)90013-Q.

Esponda, Ignacio and Pouzo, Demian. "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica*, May 2016, *84*(3), pp. 1093-30; https://doi.org/10.3982/ECTA12609.

Evans, George W. and Honkapohja, Seppo. "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations." Federal Reserve Bank of San Francisco *Economic Review*, 1993, *1*, pp. 3-13.

Evans, George W. and Honkapohja, Seppo. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press, 2001; https://doi.org/10.1515/9781400824267.

Evans, George W.; Honkapohja, Seppo; Sargent, Thomas J. and Williams, Noah. "Bayesian Model Averaging, Learning, and Model Selection," in T.J. Sargent and J. Vilmunen, eds., *Macroeconomics at the Service of Public Policy*. New York: Oxford University Press, 2013, pp. 99-119; https://doi.org/10.1093/acprof:oso/9780199666126.003.0007.

Fudenberg, Drew, and Levine, David K. "Self-Confirming Equilibrium." *Econometrica*, May 1993, *61*(3), pp. 523-45; https://doi.org/10.2307/2951716.

Geweke, John. *Complete and Incomplete Econometric Models*. Princeton: Princeton University Press, 2010; https://doi.org/10.1515/9781400835249.

Gilboa, Itzhak; Postlewaite, Andrew W. and Schmeidler, David. "Probability and Uncertainty in Economic Modeling." *Journal of Economic Perspectives*, Summer 2008, *22*(3), pp. 173-88; https://doi.org/10.1257/jep.22.3.173.

Hansen, Lars P. "Uncertainty Outside and Inside Economic Models." *Journal of Political Economy*, October 2014, *122*(5), pp. 945-87; https://doi.org/10.1086/678456.

Hansen, Lars P. and Sargent, Thomas J. *Robustness*. Princeton: Princeton University Press, 2008.

Kalai, Ehud and Lehrer, Ehud. "Rational Learning Leads to Nash Equilibrium." *Econometrica*, September 1993, *61*(5), pp. 1019-45; https://doi.org/10.2307/2951492.

Kandori, Michihiro; Mailath, George J. and Rob, Rafael. "Learning, Mutation and Long Run Equilibria in Games." *Econometrica*, January 1993, *61*(1), pp. 29-56; https://doi.org/10.2307/2951777.

Kurz, Mordecai. "On the Structure and Diversity of Rational Beliefs." *Economic Theory*, November 1994, *4*(6), pp. 877-900; https://doi.org/10.1007/BF01213817.

Lucas, Robert E. Jr. "Expectations and the Neutrality of Money." *Journal of Economic Theory*, April 1972, *4*(2), pp. 103-24; https://doi.org/10.1016/0022-0531(72)90142-1.

Perron, Pierre and Shiller, Robert J. "Testing the Random Walk Hypothesis: Power versus Frequency of Observation." *Economics Letters*, 1985, *18*, pp. 381-86; https://doi.org/10.1016/0165-1765(85)90058-8.

Sargent, Thomas J. *Bounded Rationality in Macroeconomics*. Oxford and New York: Oxford University Press and Clarendon Press, 1993.

Timmermann, Allan. "Forecast Combinations," in G. Elliott, C. Granger and A. Timmermann, eds., *Handbook of Economic Forecasting*. Volume 1. Amsterdam: Elsevier, 2006, pp. 135-196; https://doi.org/10.1016/S1574-0706(05)01004-9.

Terrorism, Trade, and Welfare

Subhayu Bandyopadhyay, Todd Sandler, and Javed Younas

For a standard competitive trade model, the authors show that the incidence of terrorism in different nations can affect the pattern of trade. Nations with a greater incidence of terrorism will export goods that are more immune to terrorism-related disruptions, while importing more terrorism-impacted goods. In addition, terrorism can be welfare augmenting for some nations because of terms-of-trade externalities. Finally, the authors present some qualitative conditions that identify when a nation's trade volume may rise (or fall) in response to a greater incidence of terrorism. Given the differential impact across nations, these trade and welfare results point to potential difficulties in international coordination of counterterrorism policy. (JEL F11, F52, H56)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2017, *99*(3), pp. 295-306.
https://doi.org/10.20955/r.2017.295-306

1 INTRODUCTION

Terrorism is the premeditated use of or threat to use violence by individuals or subnational groups to obtain a political or social objective through the intimidation of a large audience beyond that of the immediate victims (Enders and Sandler, 2012). Civil wars or internal conflicts may, but need not, involve terrorism as a tactic (Sambanis, 2008). Moreover, terrorism typically occurs in the absence of a civil war—e.g., terrorism campaigns by the Italian Red Brigade, Direct Action, and the Red Army Faction in the 1970s and 1980s. Civil wars kill many more people than terrorism, which usually involves relatively few deaths and injuries each year.[1] To be classified as a civil war, a death threshold of at least 1,000 must be reached, while there is no such threshold for terrorism. Civil wars and interstate wars are shown to have much larger economic impacts than terrorism because these wars destroy much more infrastructure, capital, and lives. Empirical studies by Blomberg, Hess, and Orphanides (2004) and Gaibulloev and Sandler (2009, 2011) indicate that internal and external wars have a much greater impact on per capita income growth and other macroeconomic aggregates for samples that include the world, Asia, and Africa.[2] In a recent study, Gaibulloev and Younas (2016)

Subhayu Bandyopadhyay is an economist and research officer at the Federal Reserve Bank of St. Louis. Todd Sandler is the Vibhooti Shukla Professor of Economics and Political Economy and director of the Center for Global Collective Action at the University of Texas at Dallas. Javed Younas is an associate professor of economics at American University of Sharjah.

find that civil wars, unlike terrorism, have a robust negative impact on domestic bank lending, which funds investment and development in developing countries.

In contrast to civil wars, terrorism is usually more sector specific, aimed at tourism, the export sector, or foreign direct investment (FDI) (Sandler and Enders, 2008). Many of these sectors also tend to be more capital intensive than other sectors so that terrorism does not impact all productive factors equally. This sector specificity of terrorism arises, among other reasons, from the publicity that terrorist organizations seek for their attacks. This publicity helps these organizations advertise their strength and recruit sympathetic individuals who share their grievances. If a terrorist organization kills two people working in a sugar plantation in India, the attack is unlikely to garner a lot of international press. On the other hand, if two American tourists perish in an attack at an Indian hotel, the terrorist incident will likely lead to a lot of press coverage—both nationally and internationally. This attention is what the terrorist organization seeks, raising the potential value to them of sectors such as tourism. Similarly, multinational firms employing citizens of different nationalities in their plants or dealerships will be relatively high-value targets. This terrorist asymmetric targeting of different productive sectors of an economy differentiates terrorism from less-targeted disruptions such as the current Syrian civil war, which makes all productive sectors vulnerable. When certain sectors are impacted more than other sectors of an economy, this targeting strategy changes the pattern of general equilibrium production and consumption. The resulting asymmetric effects on the nation's relative supply and demand of different goods (and services) will affect the nation's pattern of trade. A recent paper by Bandyopadhyay, Sandler, and Younas (2016) addresses some of these issues by building a heterogeneous-firm, monopolistic-competitive model of trade. These authors estimate the theoretical implications of their analysis by considering the effects of terrorism on trade using a sample of 151 nations. They find that terrorism tends to affect different sectors unevenly, with the greatest harm being experienced by manufacturing, compared with primary (non-manufactured) goods. Within manufacturing, they find that effects of terrorism are more damaging to high-skilled production than low-skilled production. These asymmetric effects on different sectors are also critical to the competitive model of trade that we present in this paper.

Recent high-profile terrorism attacks in France, Turkey, Belgium, Germany, and elsewhere indicate how terrorism remains a security concern that must be addressed worldwide. This threat raises questions about its economic implications, not only within a targeted nation, but also among trading partners. For the latter, terrorism may affect trade and welfare in ways that transcend the immediate loss of life and property. After the World Trade Center attacks on 9/11, terrorism-related threats have come under greater scrutiny in terms of U.S. policy. Beyond the immediate loss of life and property, there are longer-term economic consequences stemming from terrorist attacks. For example, if a country is perceived as particularly vulnerable to such attacks, foreign investors are more likely to shy away from investing in that country (Enders and Sandler, 1996); in turn, this can reduce the country's FDI, with a concomitant drop in employment opportunities. Abadie and Gardeazabal (2008), Enders, Sachsida, and Sandler (2006), and others find that terrorism can cause substantial reductions in FDI for a terrorism-afflicted nation. This FDI reduction may stem from higher risks or an FDI diversion to safer venues abroad. Terrorism can also divert government expenditures toward security activities, which may compromise other important allocations such as expenditures on a nation's infrastructure. In turn, this can affect a country's economic growth as well as its exports and imports (see Blomberg, Hess, and Orphanides, 2004; Blomberg and Hess, 2006).

Terrorism can impact international trade through various channels. For instance, it can increase the costs of transportation by raising insurance premiums for goods shipped to higher-risk nations. In turn, such higher costs augment prices of imported goods, working effectively as import tariffs to reduce trade. Terrorism can also lead to destruction of productive capacity of certain sectors of an economy, making these sectors more reliant on imports, which has the counterintuitive implication that terrorism may actually increase trade. Empirical works such as Nitsch and Schumacher (2004), Blomberg and Hess (2006), and Mirza and Verdier (2014), among others, find evidence that terrorism tends to depress trade. A recent empirical article finds that transnational terrorist attacks had virtually no short-run impact on trade.[3] This null finding suggests that the effect of terrorism on trade may be driven by factors other than transaction costs. The absence of a trade effect may be reconciled when one recognizes that terrorism can destroy capital or labor endowments unevenly in targeted sectors, thereby twisting the production possibility frontier and impacting the pattern of trade. For example, Abadie and Gardeazabal (2008) and Bandyopadhyay, Sandler, and Younas (2014) show that terrorism can reduce capital stocks by reducing FDI flows, while Bandyopadhyay and Sandler (2014b) show that such changes in factor endowments can either increase or decrease trade. The current paper differs from Bandyopadhyay and Sandler (2014b) in a few ways. First, the current paper does not use factor abundance in a Heckscher-Ohlin model to investigate terrorism-induced trade changes. Terrorism is now viewed as sector specific rather than factor specific. Second, we now abandon the small-country assumption so that there can be terrorism-induced changes to the terms of trade. Third, unlike Bandyopadhyay and Sandler (2014b), the current paper allows welfare effects that do not have to be negative in light of terrorism—i.e., there may be gainers and losers.

The existing literature, however, does not provide an analytical inquiry of international welfare implications of terrorism in a trading environment. Using a standard competitive trade model, we first identify how terrorism may affect trade patterns. Next, we show that, while some nations must lose due to increased terrorism, other nations may actually be better off due to positive terms-of-trade externalities. Finally, we provide a qualitative condition under which terrorism increases trade. Section 2 presents the model and analysis, and Section 3 contains concluding remarks.

2 A COMPETITIVE MODEL OF TRADE AND TERRORISM

Let us consider a world economy with two nations, A and B. Perfectly competitive firms in each of these nations produce homogeneous goods, x_1 (good 1) and x_2 (good 2).[4] Moreover, both of these nations may be subjected to terrorist attacks that disrupt the production process. For simplicity of exposition, we make some assumptions that are not critical to our main points. First, we assume that terrorist attacks are exogenous in the sense that we do not model their supply, nor do we consider how counterterrorism policies may reduce their impact. Second, the only impact that we consider here is on the production process.[5] Finally, we assume that the production of good 1 is adversely influenced by terrorism, while good 2 is immune to it.[6] For example, say good 1 is a manufactured good, while good 2 is a non-manufactured primary product. The manufacturing sector tends to locate in more visible urban or semi-urban areas that attract more attention from terrorist groups, thereby rendering its production process relatively more vulnerable to terrorism.[7] Both goods are produced using labor and capital, for which w and r are their respective factor prices. Firms in sector 1 use labor (L_1) and capital (K_1) to produce output x_1 with a constant returns to scale (CRS) technology, $F^1(L_1,K_1)$. A fraction $1 - \phi(T)$ of output of good 1 is lost due to terrorism-related disruptions, indexed by T, such that

(1) $$x_1 = \phi(T)F^1(L_1,K_1), 0 < \phi \leq 1, \phi'(T) < 0, \phi(0) = 1.$$

In equation (1), $\phi'(T) < 0$ implies that more terrorism results in a greater fraction of the good being destroyed.[8] Good 2 is not affected by terrorism; so, denoting labor and capital use in sector 2 by L_2 and K_2, respectively, we have

(2) $$x_2 = F^2(L_2,K_2),$$

where $F^2(L_2,K_2)$ is a standard CRS production function. We do not need to make any particular factor intensity assumptions for the purpose of this paper, except that the two sectors have different capital-to-labor ratios.

Let good 2 be the numeraire good, while the price of good 1 is p_1. The first-order conditions for profit maximization of competitive firms in the two sectors yield the familiar equalities between factor prices and the respective values of marginal products[9]:

(3a) $$p_1\phi(T)F_L^1(L_1,K_1) = w = F_L^2(L_2,K_2),$$

(3b) $$p_1\phi(T)F_K^1(L_1,K_1) = r = F_K^2(L_2,K_2).$$

These equations may be written, respectively, as

(4a) $$P_1F_L^1(L_1,K_1) = w = F_L^2(L_2,K_2), P_1 \equiv p_1\phi(T) = P_1(p_1,T),$$

(4b) $$P_1F_K^1(L_1,K_1) = r = F_K^2(L_2,K_2).$$

Equations (4a) and (4b) imply that the value of production of this economy may be represented through the standard revenue function R, which is the envelope function for the following maximization problem (noting that $p_2 = 1$)[10]:

(5) $$\text{Maximize } P_1F^1(L_1,K_1) + F^2(L_2,K_2), \text{subject to } L_1 + L_2 = \overline{L}, \text{and } K_1 + K_2 = \overline{K},$$

where $\bar{L}$ and $\bar{K}$ are the national labor and capital endowments, respectively. If we suppress factor endowments, the revenue function corresponding to equation (5) and its partial derivatives with respect to output prices are

$$\text{(6)} \qquad R = R(P_1, p_2 = 1), \text{ where } R_1(\cdot) = F^1(L_1, K_1) \text{ and } R_2(\cdot) = F^2(L_2, K_2).$$

Note from equation (6) that while $R_2(\cdot)$ yields the supply of good 2, $R_1(\cdot)$ is the supply of good 1 before the fraction $1 - \phi$ of the good is lost to terrorism-related damages. Accordingly, the actual terrorism-depleted relative supply of good 1 in terms of good 2 may be represented as

$$\text{(7)} \qquad \frac{x_1}{x_2} = \frac{\phi(T)F^1(\cdot)}{F^2(\cdot)} = \frac{\phi(T)R_1(P_1,1)}{R_2(P_1,1)} = \phi(T)\rho\left[P_1(p_1,T)\right] \equiv X(p_1,T),$$

where $\rho(P_1) \equiv \dfrac{R_1(P_1,1)}{R_2(P_1,1)}$. Standard properties like homogeneity of the revenue function of degree one in prices and also convexity of the revenue function in prices imply that $\rho'(P_1) > 0$. Using equations (4a) and (7), we have

$$\text{(8)} \qquad X_{p_1} = \phi^2 \rho'(P_1) > 0 \text{ and } X_T = \phi'(T)\left[\rho + \phi p_1 \rho'(P_1)\right] < 0.$$

Equations (7) and (8) yield a relative supply curve for good 1 that is positively sloped and shifts to the left with a greater incidence of terrorism.

We characterize the demand side by assuming a nation's preferences can be captured by a representative consumer with homothetic preferences. Consumption expenditure minimization by this consumer, at a price vector $(p_1, p_2 = 1)$ to achieve a certain utility level u, yields the consumer's expenditure function, $E(p_1, 1, u)$. Standard properties of an expenditure function yield the Hicksian demand function for good i as $E_i(p_1, 1, u)$. Given that monotonic transformations preserve ranking, there is no loss in generality if we use a homogeneous of degree one utility function for the representative consumer to denote its homothetic preferences. This expenditure function can be expressed as $E(p_1, 1, 1)u$. Accordingly, the relative Hicksian demand for good 1, RD, can be written as

$$\text{(9)} \qquad \frac{E_1(p_1,1,u)}{E_2(p_1,1,u)} = \frac{E_1(p_1,1,1)u}{E_2(p_1,1,1)u} = \frac{E_1(p_1,1,1)}{E_2(p_1,1,1)} = RD(p_1),$$

where $RD'(p_1) < 0$ follows from standard concavity properties of the expenditure function.

2.1 Effect of Terrorism on Trade

Let us assume that the two nations, A and B, are identical in all respects except for having different potential levels of terrorism. This is a reasonable assumption since nations are targeted differently owing to their foreign policy, their location, their counterterrorism actions, and their vulnerabilities (Enders and Sandler, 2006; Savun and Phillips, 2009). Let the incidence of terrorism in nation j be T^j. Given identical preferences and technology, the two nations share the same functional forms for relative demand and the relative supply of good 1.

Figure 1

Effect of Terrorism on Autarky Prices

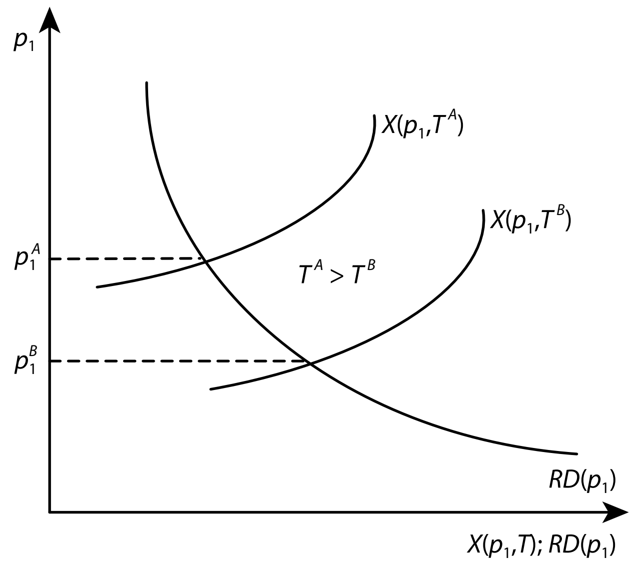

In autarky, each nation's relative demand for good 1 has to match its relative supply of good 1; hence, we have

(10a) $$X\left(p_1,T^A\right)=RD\left(p_1\right)\text{ and}$$

(10b) $$X\left(p_1,T^B\right)=RD\left(p_1\right).$$

Proposition 1

The nation experiencing more terrorism imports the terrorism-susceptible good and exports the terrorism-immune good.

Proof: If $T^A > T^B$, then $X(p_1,T^A)$ must be less than $X(p_1,T^B)$ at any p_1 (see equation (8)). This is shown in Figure 1, where the autarky price of good 1 in nation A must exceed the corresponding price in nation B. This establishes that the nation less impacted by terrorism has a comparative advantage in producing good 1, the terrorism-susceptible good. Consequently, this less-impacted nation will export good 1 and import the terrorism-immune good 2. ■

Comment: Greater incidence of terrorism in country A causes greater disruptions in the production of good 1, thus shrinking its relative supply of good 1 to a greater extent than in nation B. Given identical relative demand in the two nations, the relative scarcity of the good in nation A makes the good more expensive in A, compared with B under autarky. In turn, under free trade, the nations can enjoy gains from trade when nation A exports good 2 and imports good 1.

2.2 The Trading Equilibrium

Under free trade, the global supply of good 1 must equal the global demand for good 1 at the market-clearing equilibrium price. By Walras's law, the market for good 2 also clears at this equilibrium price. Note that the supply of good 1 in nation j is $\phi(T^j)R_1[P_1(p_1,T^j),1]$, while at a utility level u^j, its demand for good 1 is $E_1(p_1,1,u^j)$. Thus, aggregating the demand and supply in the two nations gives the market-clearing condition for good 1 under free trade as

$$\phi\left(T^A\right)R_1\left[P_1\left(p_1,T^A\right),1\right]+\phi\left(T^B\right)R_1\left[P_1\left(p_1,T^B\right),1\right]=E_1\left(p_1,1,u^A\right)+E_1\left(p_1,1,u^B\right). \tag{11}$$

Let exp_i^j be the net export of good i (i = 1,2) of nation j, which is represented by the difference between the nation's supply and demand of good i. Accordingly, we have

$$exp_1^j=\phi\left(T^j\right)R_1\left[P_1\left(p_1,T^j\right),1\right]-E_1\left(p_1,1,u^j\right) \text{ and} \tag{12a}$$

$$exp_2^j=R_2\left[P_1\left(p_1,T^j\right),1\right]-E_2\left(p_1,1,u^j\right). \tag{12b}$$

If nation j is an exporter (importer) of good i, then exp_i^j is positive (negative). Now, each nation's consumption spending must equal its production revenues.[11] Thus,

$$E\left(p_1,1,u^j\right)=R\left[P_1\left(p_1,T^j\right),1\right],\ j=A,B. \tag{13a}$$

Recalling that $E(p_1,1,u) = E(p_1,1,1)u$, we can express equation (13a) as

$$\begin{aligned} &u^j=\frac{R\left[P_1\left(p_1,T^j\right),1\right]}{E\left(p_1,1,1\right)}=u^j\left(p_1,T^j\right),\ u_{p_1}^j=\frac{exp_1^j}{E\left(p_1,1,1\right)}>0, \text{ if and only if } exp_1^j>0, \text{ and} \\ &u_{T^j}^j=\frac{p_1R_1\left[P_1\left(p_1,T^j\right),1\right]\phi'\left(T^j\right)}{E\left(p_1,1,1\right)}<0,\ j=A,B. \end{aligned} \tag{13b}$$

The term $u_{p_1}^j$ is positive if and only if nation j is an exporter of good 1 because, ceteris paribus, a rise in the price of good 1 is a terms-of-trade gain (loss) for nation j if it is an exporter (importer) of good 1. On the other hand, $u_{T^j}^j$ is unambiguously negative because either nation, ceteris paribus, loses productive capacity when the incidence of terrorism rises in that nation.

When equations (13a) and (13b) are substituted into (11), the market-clearing price of good 1 is implicitly defined as

$$p_1=p_1\left(T^A,T^B\right),\ \frac{\partial p_1}{\partial T^j}>0,\ j=A,B. \tag{14}$$

The proof for equation (14) follows standard methods and is available on request. If the international equilibrium is stable in the sense that the price of good 1 must rise whenever its global demand exceeds its global supply, then a terrorism-induced reduction in the supply of good 1 in either of the two nations must raise the world price of good 1. In turn, this implies that terrorism, regardless of its location, must improve the terms of trade of the nation exporting good 1.

Proposition 2

A rise in terrorism in either of two trading nations must reduce the welfare of the relatively terrorism-prone nation (i.e., *A*). The relatively terrorism-free nation (i.e., *B*) gains when terrorism rises in *A*, but may lose when terrorism rises in its homeland.

Proof: Using equations (13a), (13b), and (14) and noting that *A* is an importer of good 1 (i.e., $exp_1^A < 0 \Rightarrow u_{p_1}^A < 0$), we have that *A*'s welfare change, with respect to increases in terrorism in *A* and *B*, respectively, may be written as

$$\frac{\partial u^A}{\partial T^A} = u_{p_1}^A \frac{\partial p_1}{\partial T^A} + u_{T^A}^A < 0 \text{ and } \frac{\partial u^A}{\partial T^B} = u_{p_1}^A \frac{\partial p_1}{\partial T^B} < 0. \tag{15a}$$

Thus, regardless of whether terrorism rises in *A* or *B*, *A*'s welfare necessarily falls. Noting that *B* is the exporter of good 1 ($exp_1^B > 0 \Rightarrow u_{p_1}^B > 0$), we have

$$\frac{\partial u^B}{\partial T^A} = u_{p_1}^B \frac{\partial p_1}{\partial T^A} > 0 \text{ and } \frac{\partial u^B}{\partial T^B} = u_{p_1}^B \frac{\partial p_1}{\partial T^B} + u_{T^B}^B > 0 \text{ if and only if } u_{p_1}^B \frac{\partial p_1}{\partial T^B} > \left| u_{T^B}^B \right|. \tag{15b}$$

Thus, nation *B* necessarily gains when there is more terrorism in nation *A*, but it may gain or lose when it suffers more terrorism at home. ■

Comment: A rise in terrorism in nation *A* adversely affects its well-being in two ways. First, there is the direct loss in income due to terrorist attacks at given terms of trade. Second, nation *A* suffers from a rise in its import price due to a fall in the supply of the terrorism-susceptible good. Nation *B*, however, must gain when terrorism rises in *A* because it suffers no direct loss, while at the same time enjoying a terms-of-trade benefit as the price of good 1 (its export good) rises. Following similar logic, we can conclude that if, instead, *B* experiences a rise in terrorism, it will have conflicting direct and terms-of-trade effects, rendering its aggregate welfare effect ambiguous. However, even in this case, nation *A* will suffer welfare loss because its only consequence is an adverse terms-of-trade effect.

2.3 Multicountry Analysis

The previous analysis can be easily extended to a multicountry context. If nations are indexed by *A*, *B*, *C*,… the equation corresponding to equation (14) is

$$p_1 = p_1\left(T^A, T^B, T^C, \ldots\right), \frac{\partial p_1}{\partial T^j} > 0, \; j = A, B, C, \ldots \tag{16}$$

Following Proposition 1, the nation with the highest (lowest) terrorism index must be an importer (exporter) of good 1. Nations in between these two extremes may either be exporters or importers of good 1.

2.4 Small Open-Economy Equilibrium and Volume of Trade

For a sufficiently large number of relatively symmetric nations (except for some differences in their terrorism levels), we can assume that each nation is "small" in the sense that the international terms of trade are not affected by a rise in terrorism in that nation alone. In this case, results of the previous subsections continue to hold, with the caveat that $\frac{\partial p_1}{\partial T^j} = 0.$ The trade balance for nation *j* requires that the sum of the value of its net exports equals zero, so that

$$(17) \qquad exp_2^j = -p_1 exp_1^j = p_1 imp_1^j,$$

where imp_i^j is the volume of net import of good *i* for nation *j*. Given p_1, we can use either exp_2^j or imp_1^j as a measure of the volume of trade. Now, the following holds:

$$(18) \qquad exp_2^j = R_2\left[P_1\left(p_1, T^j\right), 1\right] - E_2\left[p_1, 1, u^j\left(p_1, T^j\right)\right].$$

Proposition 3
If the nation experiencing more terrorism is a net exporter of good 2, then its trade volume will increase. If, instead, it is a net importer of good 2, then its trade volume will decrease.

Proof: Differentiating equation (18) and using equations (4a) and (13b), we get

$$(19) \qquad \frac{dexp_2^j}{dT^j} = R_{21} p_1 \phi'\left(T^j\right) - E_2\left(p_1, 1, 1\right) u_{T^j}^j > 0, \text{ because } R_{21} < 0.$$ [12]

If *j* is a net exporter of good 2, $exp_2^j > 0$, then $\frac{dexp_2^j}{dT^j} > 0$ so that its exports of good 2 and, hence, its imports of good 1 must both rise. If, however, it is a net importer of good 2, then $exp_2^j = -\,imp_2^j < 0$ and $\frac{dexp_2^j}{dT^j} > 0 \Rightarrow \frac{dimp_2^j}{dT^j} < 0.$ The last inequality establishes the second part of the proposition. ■

Comment: The first term on the right-hand side of equation (19) is the increase in nation *j*'s supply of good 2 when terrorism reduces net returns in *j*'s sector 1 and causes a resource inflow to its sector 2. The second term is the reduction in *j*'s demand for good 2 through the income effect as income falls from terrorism-related damages to *j*'s productive capacity. The rise in production of good 2 and the fall in its demand must increase j's net export of good 2 when this country experiences more terrorism. Consequently, if nation *j* is an exporter of good 2, terrorism raises its exports and the volume of trade rises. But, if it is an importer of good 2, the rise in good 2's production and fall in the good's demand must reduce import demand. In this case, terrorism reduces *j*'s trade volume.

3 CONCLUDING REMARKS

The paper establishes a Heckscher-Ohlin-type result that nations that are more susceptible to terrorism are likely to import goods that are more subject to terrorism-related disruptions and to export other goods. This insight also suggests that there is no reason to believe that terrorism necessarily reduces trade. While higher trading costs tend to reduce trade, changes in production patterns, as well as incomes of nations, can lead to general-equilibrium reallocations that can raise trade. Therefore, whether terrorism reduces or raises trade is a context-specific issue.

Our welfare analysis shows that, while terrorism is necessarily welfare reducing for a nation that imports the greater terrorism-impacted good, the exporting nation may or may not be worse off because of potential gains from trade. If, indeed, a nation does not directly suffer from terrorist attacks, then it must gain (lose) due to terms-of-trade effects of terrorism if it is an exporter (importer) of the terrorism-disrupted good. In turn, this may make international coordination in counterterrorism policies more difficult to achieve because nations that gain from terrorism will have no incentive to participate in such an international coalition. This is discouraging because international coordination has been difficult to achieve under the best of circumstances as nations are reluctant to coordinate on security measures that could jeopardize their autonomy (Enders and Sandler, 2012, Chapter 6). ■

NOTES

1 Transnational terrorist incidents are associated with just over one death per incident on average, and there are from 200 to 500 such incidents a year (Enders and Sandler, 2012). Of course, the four hijackings on September 11, 2001 (henceforth 9/11) is an exception with close to 3,000 fatalities. The relatively few high-profile incidents in the recent news give a misleading impression of death and destruction.

2 These studies directly compare the economic impacts of wars and terrorism. To our knowledge, there are no studies that contrast the economic effect of terrorism with that of crime, the latter of which is not politically motivated. Obviously, crime is far more prevalent in societies than terrorism. Crime is often motivated by personal economic gain, and in these cases will be opportunistically targeted at less-protected individuals or property for maximal economic benefits. Terrorism is usually motivated by ideologies and targeted at nations, authorities, and individuals opposed to that ideology. So, if the symbolic or advertisement value of a terrorist attack is large, it may be optimal to engage in it, even when there are small odds of success and only modest economic damage is inflicted. Accordingly, airports or public buildings are more vulnerable to terrorism, while less-protected urban areas in the United States are more vulnerable to crime.

3 See Egger and Gassebner (2015), who find no terrorism impact within 1.5 years of an attack.

4 Until Section 2.1, we do not need to distinguish between the two nations because the same model applies to both.

5 Terrorism is a multi-faceted issue that may affect both producers and consumers. Its incidence is also generally endogenous to a nation's policymakers through proactive counterterrorism policies against the perpetrators or through defensive counterterrorism policies. The latter policies reduce the damage from attacks or deflect such attacks abroad. In economics and political science, there is a large and emerging literature on this topic. See Enders and Sandler (2012), who offer a comprehensive discussion of the literature. For an explicit treatment of counterterrorism policies and their impact on decisionmaking by a terrorist organization, see, e.g., Bandyopadhyay and Sandler (2014a).

6 This assumption is not critical and can be easily relaxed. All that is required to establish the results that follow is that sector 1 is more vulnerable to terrorism compared with sector 2, such that the supply of good 1 relative to the supply of good 2 is negatively impacted by terrorism.

7 As discussed earlier, terrorists typically tend to target more-prominent urban areas to achieve maximum publicity for their attacks. Relatively urban and industrial sectors are higher-value targets compared with rural agricultural areas. Indeed, Bandyopadhyay, Sandler, and Younas (2016) find that terrorism has more severe negative effects on trade in manufactured goods compared with primary products.

8 This is a particular conceptualization of how terrorism-related production disruptions may be manifested, where one of the two goods is more susceptible to such disruptions. One interpretation is that factors are used to produce a unit of output, part of which is damaged or lost before being sold to a consumer. An alternative characterization is that labor and capital become proportionally less productive due to terrorism, such that each unit of labor and capital effectively becomes $\phi(<1)$ units. Because of constant returns to scale, this yields the same production function as described in equation (1).

[9] When we use subscripts in function notations, they refer to partial derivatives with respect to the respective argument of the function.

[10] For details on the properties of revenue functions used in trade theory, see Dixit and Norman (1980).

[11] This is equivalent to assuming trade is balanced in the sense that, at the world equilibrium price, the value of exports of both nations equals the value of their imports.

[12] Convexity and homogeneity properties of the revenue function ensure that $R_{21} < 0$.

REFERENCES

Abadie, Alberto and Gardeazabal, Javier. "Terrorism and the World Economy." *European Economic Review*, January 2008, *52*(1), pp. 1-27; https://doi.org/10.1016/j.euroecorev.2007.08.005.

Bandyopadhyay, Subhayu and Sandler, Todd. "Immigration Policy and Counterterrorism." *Journal of Public Economics*, February 2014a, *110*, pp. 112-23; https://doi.org/10.1016/j.jpubeco.2013.12.013.

Bandyopadhyay, Subhayu and Sandler, Todd. "The Effects of Terrorism on Trade: A Factor Supply Approach." Federal Reserve Bank of St. Louis *Review*, Third Quarter 2014b, *96*(3), pp. 229-41; https://research.stlouisfed.org/publications/review/2014/q3/bandyopadhyay.pdf.

Bandyopadhyay, Subhayu; Sandler, Todd and Younas, Javed. "Foreign Direct Investment, Aid, and Terrorism." *Oxford Economic Papers*, 2014, *66*(1), pp. 25-50; https://doi.org/10.1093/oep/gpt026.

Bandyopadhyay, Subhayu; Sandler, Todd and Younas, Javed. "Trade and Terrorism: A Disaggregated Approach." Working Paper No. 2016-001A, Federal Reserve Bank of St. Louis, February 2016; https://research.stlouisfed.org/wp/2016/2016-001.pdf.

Blomberg, Stephen B.; Hess, Gregory D. and Orphanides, Athanasios. "The Macroeconomic Consequences of Terrorism." *Journal of Monetary Economics*, July 2004, *51*(5), pp. 1007-32; https://doi.org/10.1016/j.jmoneco.2004.04.001.

Blomberg, Stephen B. and Hess, Gregory D. "How Much Does Violence Tax Trade?" *Review of Economics and Statistics*, November 2006, *88*(4), pp. 599-612; https://doi.org/10.1162/rest.88.4.599.

Dixit, Avinash K. and Norman, Victor. *Theory of International Trade: A Dual, General Equilibrium Approach*. Cambridge: Cambridge University Press, 1980; https://doi.org/10.1017/CBO9780511628627.

Egger, Peter and Gassebner, Martin. "International Terrorism as a Trade Impediment?" *Oxford Economic Papers*, January 2015, *67*(1), pp. 42-62; https://doi.org/10.1093/oep/gpu037.

Enders, Walter; Sachsida, Adolfo and Sandler, Todd. "The Impact of Transnational Terrorism on U.S. Foreign Direct Investment." *Political Research Quarterly*, December 2006, *59*(4), pp. 517-31; https://doi.org/10.1177/106591290605900402.

Enders, Walter and Sandler, Todd. "Terrorism and Foreign Direct Investment in Spain and Greece." *Kyklos*, August 1996, *49*(3), pp. 331-52; https://doi.org/10.1111/j.1467-6435.1996.tb01400.x.

Enders, Walter and Sandler, Todd. "Distribution of Transnational Terrorism Among Countries by Income Classes and Geography After 9/11." *International Studies Quarterly*, June 2006, *50*(2), pp. 367-93; https://doi.org/10.1111/j.1468-2478.2006.00406.x.

Enders, Walter and Sandler, Todd. *The Political Economy of Terrorism*. 2nd Edition. New York: Cambridge University Press, 2012.

Gaibulloev, Khusrav and Sandler, Todd. "The Impact of Terrorism and Conflicts on Growth in Asia." *Economics & Politics*, November 2009, *21*(3), pp. 359-83; https://doi.org/10.1111/j.1468-0343.2009.00347.x.

Gaibulloev, Khusrav and Sandler, Todd. "The Adverse Effect of Transnational and Domestic Terrorism on Growth in Africa." *Journal of Peace Research*, May 2011, *48*(3), pp. 355-71; https://doi.org/10.1177/0022343310395798.

Gaibulloev, Khusrav and Younas, Javed. "Conflicts and Domestic Bank Lending." *Public Choice*, December 2016; *169*(3), pp. 315-331; https://doi.org/10.1007/s11127-016-0362-3.

Mirza, Daniel and Verdier, Thierry. "Are Lives a Substitute for Livelihoods? Terrorism, Security, and U.S. Bilateral Imports." *Journal of Conflict Resolution*, September 2014, *58*(6), pp. 943-75; https://doi.org/10.1177/0022002713487312.

Nitsch, Volker and Schumacher, Dieter. "Terrorism and International Trade: An Empirical Investigation." *European Journal of Political Economy*, June 2004, *20*(2), pp. 423-33; https://doi.org/10.1016/j.ejpoleco.2003.12.009.

Sambanis, Nicholas. "Terrorism and Civil War," in P. Keefer and N. Loayza, eds., *Terrorism, Economic Development, and Political Openness*, Cambridge: Cambridge University Press, 2008, pp. 174-206; https://doi.org/10.1017/CBO9780511754388.007.

Sandler, Todd and Enders, Walter. "Economic Consequences of Terrorism in Developed and Developing Countries," in P. Keefer and N. Loayza, eds., *Terrorism, Economic Development, and Political Openness*. Cambridge: Cambridge University Press, 2008, pp. 17-47; https://doi.org/10.1017/CBO9780511754388.002.

Savun, Burcu and Phillips, Brian J. "Democracy, Foreign Policy, and Terrorism." *Journal of Conflict Resolution*, December 2009, *53*(6), pp. 878-904; https://doi.org/10.1177/0022002709342978.