

MOVING TO JOBS: THE ROLE OF INFORMATION IN MIGRATION DECISIONS*

Riley Wilson

July 10, 2020

Click [here](#) for latest version

Abstract

This paper exploits county-level variation in exposure to news about labor markets impacted by fracking to show that access to information about employment opportunities affects migration. Exposure to newspaper articles about fracking increased migration to areas mentioned in the news by 2.4 percent on average, concentrated among young, unmarried, less-educated men. Commuting also increased, sentiment of the news matters, and TV news has an impact. Google searches for “fracking” and the names of states specifically mentioned spike after news broadcasts about fracking. Counties experiencing weak labor markets are the most responsive, suggesting these areas see large benefits to information provision.

Keywords: geographic mobility, migration, information, news, fracking

JEL Codes: J61, D83, R23, Q33

*Brigham Young University, Department of Economics, 435 Crabtree Building, Provo UT, 84602. Email: riley_wilson@byu.edu. I am grateful to Melissa Kearney, Judy Hellerstein, Katharine Abraham, Sergio Urzua, Ethan Kaplan, Nolan Pope, Joe Price, Cody Tuttle, Fernando Saltiel, as well as individuals at the University of Maryland, Brigham Young University, NC State, University of Arizona, University of Utah Finance Department, Federal Reserve Board, Atlanta Federal Reserve Bank, New York Federal Reserve Bank, Upjohn Institute, Southern Economic Association, the Chicago Labor Dynamism Conference, and SOLE for useful comments. I gratefully acknowledge funding through the Economic Club of Washington Vernon E. Jordan Jr. doctoral fellowship.

I Introduction

Migration is often seen as a way for people to avoid weak labor markets, and encounter better economic opportunities (Shultz, 1961; Sjastaad, 1962). However, empirically people are unlikely to move away when labor markets do poorly, even when labor markets with better employment prospects exist (Monras, 2015). As the previous literature notes, low migration might be the outcome of optimal decision-making, but might also be the result of market frictions such as credit constraints or incomplete information.

While incomplete information can introduce uncertainty and increase the risk associated with migration, information provision can increase the perceived return to moving and change migration behavior. However, the link between information and migration is frequently overlooked in empirical work. In this paper I evaluate the role of information in migration decisions by asking, does information in the news about potential, lucrative employment opportunities in other labor markets induce people to move to those markets? I am able to evaluate this relationship by exploiting the unique setting created by the fracking boom.

Over the last 10 to 15 years, the combination of horizontal drilling and hydraulic fracturing has led to localized “fracking booms” and sudden, large increases in both local employment and earnings. These booms have created large, persistent effects across industries (Feyrer, Mansur, & Sacerdote, 2017; Maniloff & Mastromonaco, 2014), resulting in perceived net benefits (Bartik et al., 2017). This gold rush-like flurry of economic activity has led to numerous newspaper articles and television news broadcasts touting the local economic impacts of fracking or debating its adverse side effects. As such, fracking has created plausibly exogenous labor market improvements that get talked about in the news. The novelty of fracking also introduced new words and vocabulary, making it easy to parse newspaper and TV news to see which sources are talking about fracking, which places they are talking about, what they are saying, and when they are saying it. Combining content about specific fracking destinations with origin-level measures of news circulation, allows me to isolate one

particular source of information transmission and estimate its causal impact on migration.

To understand the identifying variation I exploit, consider the following example. In 2012 the *USA TODAY* published six articles about fracking in Pennsylvania. My strategy tests to see if origins that had historically higher circulation of the *USA TODAY* (and thus higher exposure) saw larger increases in migration to Pennsylvania fracking counties when this news was distributed, relative to counties with historically lower circulation. In essence, this specification holds fixed any characteristic of the destination that might be changing over time, and relies on variation across origins in historic circulation to identify the effect of news exposure on migration.

I generalize this in a regression framework by combining all of the news articles about fracking from major national newspapers, with proprietary pre-fracking county-level circulation data. I also include origin by destination fixed effects and destination by year fixed effects. The destination by year fixed effects control for characteristics of the destination that are changing over time and make this a comparison of migration flows from different origins, with different levels of circulation exposure, to the same destination. The origin by destination fixed effects controls for time invariant pair-specific characteristic (such as distance) that might affect migration behavior, but it also controls for the fact that people in origins with higher circulation might be more educated, more wealthy, or more mobile on average.

I find that exposure to national newspaper news about a particular destination state increases the flow of migrants to fracking counties in that state. During my analysis period the average impact of exposure to fracking news is a 2.4 percent increase in annual origin-destination specific migration flows. Exposure to news about fracking also increases cross county commute flows by 6.6 percent on average, translating into an average annual increase of approximately one to two additional migrants and four additional commuters to the fracking destination from each origin. Although this response is small, it is economically significant given the scope of the “treatment” and the aggregate effects at the destination.

In 2012 alone, the national news about local fracking booms increased migration flows to fracking counties by 4.2 percent on average, and increased commute flows by 11.7 percent. As information comes from other source as well, this likely represents a lower bound of the overall effect of information on migration. The migration responses are largest among men, young workers (under 34), the unmarried, and workers with some college but no degree.

One concern with this strategy is that the pre-fracking level differences in circulation – which generate the identifying variation – might be correlated with other origin level characteristics that are changing over time and affect migration. For example, counties with high readership of the *USA TODAY* might be more affluent and see more income growth over time. If this additional growth in income affects migration decisions, the estimates would be biased. However this does not appear to drive the results. Areas with high and low historic circulation follow similar trends in migration to fracking areas in the pre-period, and only diverge once they are “treated” with the news. The estimates are also insensitive to controlling for time-varying origin level characteristics, like average earnings or unemployment. Furthermore, because there are 16 states involved in fracking, I can include origin by year fixed effects and account for any observable or unobservable characteristics of the origin that are changing over time and affect migration. This controls for the possibility that counties with high circulation might be becoming more educated, wealthy, or dynamic over time, in ways that might affect migration decisions. The estimates are unchanged when including the origin by year fixed effects and robust to controls for local news exposure, various functional forms, sample restrictions, and an alternative strategy comparing neighboring counties on either side of a local newspaper’s distribution market.

Given the robustness of this result, I conduct additional analyses to better understand how the news influences migration behavior. News about fracking in a particular state increases migration to that state but not other fracking states, suggesting the news conveys a location signal. Similarly, the effect of news exposure varies with distance, peaking for counties 400 to 1,000 miles away from the potential fracking destination, consistent with

people already being aware of nearby opportunities, but lacking information about distant labor market opportunities. “Positive” labor market news articles that discussing things like, jobs, booms, or growth, have a larger positive effect on migration than “negative” environmental news articles, discussing contamination, pollution or earthquakes. The effect of negative environmental news is still positive, suggesting even this news might provide information about where fracking is occurring. Positive and negative news affect cross-county commuting similarly, consistent with non-resident workers mostly experience the benefits of fracking while not incurring many of the costs (e.g., potential home water contamination). Using an analogous strategy that looks simultaneously at both newspaper exposure and TV news exposure, I find responses to both, suggesting other information sources matter as well. To understand how this information is used, I show that in the days following a news broadcast about fracking, Google search interest in the term “fracking” and the names of specific states mentioned in the news spike, consistent with people going online to seek more information. There are similar spikes in fracking related tweets.

The data suggest that the effect of newspaper exposure is over twice as large in origin counties with weak labor markets as it is in stronger labor markets, even though they face similar levels of exposure. This would suggest that all else equal, providing labor market information can be a way of increasing geographic mobility, and might be particularly effective if targeted toward weak labor markets where the returns to migration are plausibly the largest and where we have also observed non-responsiveness in the past. As this strategy has only focused on a few sources of information, the overall impact of information provision on migration behavior is potentially much larger.

II Information in Migration Decisions

There is a large literature exploring the migration response to local labor market conditions, and documenting how this response varies by demographics, educational attainment, and geography (Bound & Holzer, 2000; Wozniak, 2010; Molloy, Smith,& Wozniak, 2011).

Although heterogeneity in preferences or costs might contribute to these differences (Notowidigdo, 2013; Ganong & Shaog, 2017), there is credible evidence that liquidity constraints, credit constraints, and other market frictions impact the migration decision (Kling, Liebman, & Katz, 2007; Bryan, Chowdhury, & Mobarak, 2014). One potential friction is a lack of information.

The theoretical work has long recognized that information will affect migration decisions, but the empirical work has largely been limited to focusing on the role of networks or linguistic and cultural enclaves.¹ There are a few exceptions. Results from the Moving to Opportunity (MTO) experiment and related work suggest that providing guidelines and information about local neighborhood poverty levels along with housing vouchers and assistance induced households to move to more affluent neighborhoods (Kling et al., 2007; Bergman et al., 2019). Although the MTO did not improve economic outcomes for treated adults, it did have positive long-run impacts on young children (Chetty, Hendren, & Katz, 2015). Exploiting the Vietnam draft, Malamud and Wozniak (2012) show that college attendance increased migration rates, plausibly by increasing information about other areas through peer exposure.² Using a publicized rating system of social service offices in the UK, McCauley (2019) shows that information provision induces welfare migration. Kaplan and Schulhofer-Wohl (2017) propose a structural framework where information helps people learn about amenities in a different location. A similar information updating process can be applied to people’s expectations about labor market opportunities to see how labor market

¹See for example Greenwood (1975), Winters, de Janvry, & Sadoulet (2001), Munshi (2003), McKenzie & Rapoport (2007, 2010), and Hanson & McIntosh (2010).

²There is also work in the developing context suggesting that provision of labor market information in Bangladesh only impacts migration when combined with a conditional cash transfer (Bryan et al., 2014) and that access to more TV stations in Indonesia reduced the likelihood of moving (plausibly by correcting overly optimistic expectations about the returns to migration (Farre & Fasani, 2013). It is difficult to generalize these results to the United States. For example, the conditional round-trip transfer in Bangladesh was only equal to \$8.50 (about one weeks work), suggesting these people are highly credit constrained (Bryan et al., 2014).

information might impact migration decisions.³

In the canonical migration choice model (Sjaastad, 1962), an individual will move if the lifetime utility derived from moving to destination d minus the fixed costs of moving exceeds the utility of staying at the original location (o). However, individuals likely face incomplete information about the return to moving to destination d . This lack of information could impact the individual's willingness to move (please see Appendix C for the complete conceptual model motivating this discussion). Receiving positive information about d could lead an individual to update her beliefs about the potential return to moving there, increasing her willingness to move there.

For example, individuals exposed to numerous newspaper articles or TV news broadcasts touting the local economic benefits of fracking in Texas might adjust their beliefs about average wages or employment prospects in a Texas fracking county. Getting more information about a destination d could lead individuals to update their beliefs about the return to moving to d . Even news about negative aspects of fracking (e.g., water contamination risk) can provide information about where fracking is occurring and lead people to update their beliefs.⁴ This news does not necessarily need to be correct, as long as the individual believes it contains truthful information. If information in the news does affect the migration decision, we would expect migration to increase when people's exposure to the news increases. It is also possible that each additional piece of information will have a smaller impact on migration as people become more confident in their prior.

³An early related literature explores how things like the risk of unemployment (Todaro, 1969) and uncertainty about the future affect migration and human capital investments more generally (see Becker, 1962; Greenwood 1975, 1985; Langley, 1974; O'Connell, 1997). Under uncertainty, different states of the world occur with some known probability. Under incomplete information, potential destinations, possible states of the world, and the true probabilities are potentially unobserved.

⁴Up through 2012, the last year of my sample, about 60 percent of adults were familiar with fracking, and over half of this population was in favor of fracking (Pew Research, 2013a). For someone that views fracking favorably, even a negative news story could provide information about where fracking is occurring, and result in updated beliefs.

In fact, there is evidence that exposure to news about fracking during this time increased the likelihood a person is aware of fracking and approves of it. The Pew Research Center’s March 2012 Political Survey asked respondents if they have heard “a lot”, “a little” or “nothing at all” about fracking. They then asked those who had heard anything about fracking if they approve of fracking. Using a measure of news exposure (described below) and respondents’ state of residence, I find that state-level exposure to news about fracking is associated with increases in the probability of hearing a lot about fracking and increases in the probability of approving of fracking conditional on hearing about it (see Appendix Table A1). The impacts on awareness are observed across most demographic groups, while the impacts on approval are concentrated among men, younger individuals, and individuals with some college but no degree, groups that were likely to migrate to fracking (Wilson, 2020). This would suggest that news about fracking has the potential to shape people’s beliefs, which could impact migration behavior. I will test these patterns empirically to determine how exposure to information in the news about fracking impacts migration to fracking areas.

III Setting and Data Sources

Fracking provides a unique setting to explore the impact of news exposure on migration outcomes. Fracking began quite suddenly in the mid-2000s and by 2012 had affected oil and gas production in 252 counties in 16 states. These local fracking booms increased economic activity and improved labor markets in those counties. An additional one million dollars of production value increases county-level total wages by approximately \$80,000 and increase earnings in the commuting zone by approximately \$114,000 (Feyrer et al., 2017). Fracking also created more jobs; an additional one million dollars of production value increased county-level employment rates by 0.85 percentage points. This is due to increases in mining (0.29), transportation (0.24), construction (0.12), and the government (0.10). Not only does fracking create jobs directly in oil and gas extraction, there are large, positive spillovers on employment in transportation and construction. There is also evidence of increased earnings

in education, health, and other services. This is relevant as it is not necessary for potential migrants to work directly in oil and gas extraction. Kearney and Wilson (2018), Wilson (2020), and Cascio and Narayan (2019) show that the increases in earnings and employment are largest among men without a college degree (some college or less). Local fracking booms are credited with creating as many as 640,000 new jobs (Feyrer et al., 2017).

States in all four census regions have been affected and many people were unaware of exactly where these fracking booms were occurring. Both positive and negative aspects of fracking have been highly publicized through newspapers and TV news, and many of these news stories reference specific locations affected by fracking. Because fracking is a novel term, I am able to parse news content to identify which sources discuss fracking, which places they talk about, and what aspects of fracking were discussed. By linking this with measures of news penetration, I am able to estimate how geographic differences in exposure to news about fracking affect migration flows. This estimation requires detailed data on migration flows, news content, and news circulation. In this section I briefly describe each data source with a full description in the online data appendix (Appendix B).

First, I use well-level production data from DrillingInfo to identify fracking counties in the US. A fracking county is defined as any county with positive oil or gas extraction from a non-vertical well in a drilling formation that corresponds to a shale play. There are fracking counties in 16 states: Arkansas, California, Colorado, Louisiana, Michigan, Mississippi, Montana, New Mexico, North Dakota, Ohio, Oklahoma, Pennsylvania, Texas, Utah, West Virginia, and Wyoming as seen in Figure 1. I will examine migration pattern to fracking counties in these states.

Migration Data. Migration data is obtained from the Internal Revenue Service (IRS) Statistics of Income (SOI). Using tax documentation, such as Tax Form 1040, the IRS tracks the number of households that filed their taxes in one county in one year and in a different county the next year. They then report the number of returns (households) and

tax exemptions (individuals) that move for each county pair.⁵ County pairs are censored for privacy purposes when there are fewer than 10 returns that move. In 2013, the censoring threshold increased from 10 to 20 returns, which would have suppressed 47.5 percent of the county-to-county flows between 2000 and 2012. For this reason I restrict my analysis to migration between 2000 and 2012.⁶ I then aggregate up flows from each origin county to fracking counties in each fracking state. Because the IRS only provides a total count, I will use 2005 to 2012 American Community Survey microdata obtained from IPUMS to explore heterogeneous effects by demographic characteristics (Ruggles et al., 2015). This allows me to see if exposure to news about fracking affects migration differently by gender, race, education, and marital status.

Newspaper Circulation Data. Proprietary newspaper readership data is obtained from the Alliance for Audited Media (AAM). The AAM conducts regular newspaper circulation audits for national, regional, and most local newspapers in the United States. This includes the number of copies sold on the audit date and the number of copies as a percent of households for each county with over 25 copies. Counties with fewer than 25 copies sold are assigned a zero value. For some newspapers, these measures are only available at the Designated Market Area (DMA) level. Historic circulation rates from 2005 through 2008 are scraped from pdf files.

Newspaper Content Data. Newspaper content is obtained through the LexisNexis database, which provides access to articles from over 2,600 news sources, including *USA TODAY*, the *New York Times*, and the *Wall Street Journal*. First, I preserve all US based articles since

⁵In 2011, the IRS extending the data collect period from September to the end of the year. As such, households that file later (typically wealthier households) begin to be picked up in the data. The IRS have evaluated this change, and report that it had small impacts on state-level net migration rates (Pierce, 2015). If I exclude the states that experienced the largest change as reported by the IRS, the results are unchanged. See the data appendix for more detail.

⁶To the extent possible I extend this through 2015. The impacts are smaller in magnitude, due to the increased suppression, but equally significant and follow the same pattern. The patterns also hold if I only include data from 2009-2012 or 2011-2012 and omit the pre-treatment years (See Appendix Table A8).

1999 that include any of the search terms “frack~”, “fracing”, or “hydraulic fractur~” anywhere in the text. I then parse each article to exclude spurious keyword references such as “frick and frack” or people’s last names. Most of my analysis is restricted to three national news sources: *USA TODAY*, the *New York Times*, and the *Wall Street Journal*. In depth news coverage of fracking begins in 2009, and dramatically increases each year. In these three newspapers there were 562 news articles related to fracking between 1999 and 2012. The first two articles in the national news were in 2002 and 2003 in the *New York Times*, which briefly reference court cases about patents related to hydraulic fracturing. There was then one article in 2006, five in 2008, 20 in 2009, 48 in 2010, 198 in 2011, and 288 in 2012, meaning that most of the variation in content comes from the last two years of the sample. Next, I parse the entire text of each of these articles to determine which of the 16 fracking states listed above each article discusses.⁷ I report the number of articles by state in Appendix Table A2.

TV Viewership Data. TV viewership data is calculated from the 2008 Television and Cable Factbook using Nielsen viewership data. Between 2007 and 2009, TV stations were transitioning from analog to digitally transmitted broadcasts on a market-by-market basis. When a market transitioned, viewers were required to obtain digital reception equipment. This might have induced some viewers to substitute to other outlets (i.e., cable), meaning viewership rates in 2008 might be less correlated with viewership rates at the time of the broadcasts for markets that transitioned after 2008.⁸ For this reason I also examine the most recent viewership rates from 2016. TV viewership is reported at the DMA level for each TV station and is not program specific. The viewership rate is constructed by dividing total weekly viewership by the total number of households in the DMA.

⁷Not every article mentions a specific state. I have also parsed each article for city names from the U.S. Postal Service’s registry of city names, but find that local jurisdictions are referenced far less frequently.

⁸A special thanks to Matt Long from Warren Communication News for finding out how the viewership rates for the 2008 Factbook were constructed, and to Colin Wick for transcribing viewership rates from the 2008 Factbook.

TV News Content Data. TV news content is obtained from the Vanderbilt Television News Archive (VTNA). The VTNA database contains TV news recordings and transcript abstracts for nightly news broadcasts from the three major news networks (ABC, CBS, and NBC). VTNA only provides content for one hour of programming for the cable news outlets CNN and Fox News. Because cable news has limited content available and does not have reported viewership rates I restrict the sample to the three major news networks. I parse the transcript abstracts for search terms such as “fracking” and “shale” as well as which state is being discussed.⁹ Between 1999 and 2012 there is far less coverage of fracking on the nightly news than in the newspaper. The VTNA database only records 17 news broadcasts, with one in 2006, two in 2008, three in 2010, four in 2011, and seven in 2012.

Cross-County Commute Data. I also explore impacts on workers who live in one county but work in another using the Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES). This captures both long distance commuting and temporary relocation, such as moving to the job site for several weeks at a time but maintain the same permanent address. This data provides the number of jobs for each home and work census block pair which I aggregate up to the county level to measure the number of workers who live in one county but work in another. This data is available beginning in 2002, and also provides statistics by broad age groups (under 30, 30-54, over 54), monthly earnings in the job you are commuting for (under \$1,250, \$1,250-3,333, over \$3,333), and industry of the job you are commuting for (goods, trade/transportation, other). This allows me to explore heterogeneous commute responses across different groups.

⁹These search terms differ from those used in the newspaper analysis. When I restrict my search to the term “fracking”, only 12 news broadcasts during my sample window are found. Expanding to the other search terms used in the newspaper analysis (“frack~”, “fracing”, and “hydraulic fractur~”) does not add additional broadcasts. However, given that only abbreviated transcripts are provided, it is possible I am missing fracking related broadcasts. For this reason I expand my search terms to include “shale”. This adds five more unique broadcasts. The term “shale” could reference other topics, but since there are so few broadcasts to check in the TV analysis I am able to verify that they are references to fracking in a way that is not feasible with the newspaper analysis where there are thousands of articles.

County Characteristics Data. County level economic measures are obtained from the Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages (QCEW). County-level age and racial demographics are obtained from the National Cancer Institute’s Surveillance, Epidemiology, and End Results (SEER) Program. Other county level characteristics are obtained from the 2000 Census and ACS through the American Factfinder.

IV Empirical Strategy

Consider the hypothetical relationship between news exposure and Y_{oSt} , a measure of migration flows from origin county o to fracking counties in destination state S in year t :

$$Y_{oSt} = f(\text{news exposure}_{oSt}) + \phi_{oS} + \psi_{St} + \lambda_{ot} + \nu_{oSt}. \quad (1)$$

Migration flows are potentially impacted by exposure to the news, but also time invariant origin/destination specific characteristics, such as distance or industry ties; time-varying destination specific (ψ_{St}) and origin specific (λ_{ot}) characteristics, such as local labor market performance; and idiosyncratic origin-by-destination pair specific shocks (ν_{oSt}). This is potentially problematic for causal identification, as $\text{news exposure}_{oSt}$ might be correlated with circumstances at either the origin or destination.

These concerns are overcome in the following thought experiment: suppose we could randomly assign county-level exposure (i.e., circulation rates) to a newspaper that publishes news about the labor market impacts of fracking in Texas (or any other fracking state). By random assignment, news exposure will be uncorrelated with unobserved time-varying origin characteristics (λ_{ot}) and the origin/destination idiosyncratic term (ν_{oSt}). By comparing migration flows to Texas from counties with high and low circulation rates, everything about fracking counties in the destination state is held constant, allowing the effect of news exposure on migration to be isolated. If news about fracking in several states is being published, each

of the state experiments could be stacked and estimated as follows

$$Y_{ost} = f(\text{news exposure}_{ost}) + \phi_{oS} + \psi_{St} + \varepsilon_{ost}. \quad (2)$$

Origin/destination pair fixed effects (ϕ_{oS}) control for time invariant characteristics of the pair that affect migration, like distance. Destination state-by-year fixed effect (ψ_{St}) control for destination specific characteristics that are changing over time, and makes this a comparison of migration flows to the same destination state from origin counties that have different levels of news exposure. Importantly, this fixed effect captures destination-level changes in fracking production, labor market characteristics, and amenities which might directly affect migration behavior and lead to higher news exposure.

In reality, county-level exposure to the news is not randomly assigned. If people from Franklin County, Ohio start moving to Alleghany County, Pennsylvania, the local *Columbus Dispatch* might produce more content about fracking in Pennsylvania, raising concerns about reverse causality. While the content decisions of local newspapers have been shown to respond to local consumer preferences, this is less true for large national newspapers, such as *USA TODAY*, the *New York Times*, and the *Wall Street Journal*, which do not have a well-defined geographic market and operate differently (Gentzkow & Shapiro, 2010).¹⁰ Counties across the country are exposed to the same national news, regardless of how their preferences deviate from the national trend.¹¹ Exposure to this news will vary based on local readership and

¹⁰Gentzkow and Shapiro (2010) also list the *Christian Science Monitor* as a national newspaper, however, circulation for this newspaper is only available at the state level, so it is excluded from all analysis.

¹¹Although national newspapers might report more about destinations that see large changes in labor markets or migration (nationwide trends), the destination by year fixed effects compares migration flows from different origins to the same destination, eliminating destination specific differences that might drive news coverage. It could be argued that readers in and around New York City have a large effect on the content decisions of the *New York Times*. As a precaution, I exclude counties in the New York City DMA from the analysis. In Column (1) of Appendix Table A7 I show that the migration response is still significant if the New York City DMA is included.

circulation rates. Since fracking began quite suddenly, it is possible to isolate pre-existing variation in exposure that is not endogenous to preferences toward fracking. Oil and gas production due to fracking only began in earnest in 2008, with little national news attention prior to 2009.

To isolate plausibly exogenous variation in news exposure, I will focus on exposure to news from national newspapers as follows

$$newspaper\ exposure_{ost} = \sum_{n \in N} \left(\text{total articles in } n \text{ about fracking in } S \right)_t * Pre09\ circ.\ rate_{on}. \quad (3)$$

$Newspaper\ exposure_{ost}$ is the weighted sum of national news articles that mention fracking in destination state S in year t , where each newspaper is weighted by its fixed pre-2009 circulation rate (ranging from zero to one) in the origin county o .¹² N is the set of national newspapers: *USA TODAY*, *New York Times*, and the *Wall Street Journal*. $Newspaper\ exposure_{ost}$ is increasing in the number of articles about a particular destination, and increases by relatively more for counties that had high pre-fracking readership of the publishing newspaper. This strategy is similar to previous work using variation in circulation exposure to explore the impact of media and news on other outcomes.¹³

As in the thought experiment, the effect of national news exposure is identified by variation across origin counties in pre-fracking circulation rates. This is potentially problematic if pre-fracking circulation is correlated with changes over time in other local characteristics that affect preferences to move to fracking, captured in λ_{ot} . There are no obvious, strong geographic correlations in pre-2009 circulation of the *USA TODAY* and there is significant variation even among neighboring counties (Figure 2). Counties with low and high circulation of the *USA TODAY* also appear similar

¹²This is average circulation between 2005 and 2008.

¹³For example, Gentzkow (2006) examines TV introduction on voter turnout, DellaVigna and Kaplan (2007) examine Fox News introduction on Republican vote shares, Jensen and Oster (2009) examine Indian cable introduction on women's status, Chong and La Ferrara (2009) and La Ferrara et al. (2012) examine Brazilian soap opera introduction on divorce and fertility, Garthwaite and Moore (2012) examine exposure to Oprah Winfrey content on votes for Barack Obama after her endorsement, Kearney and Levine (2015a; 2015b) examine exposure to the MTV series "16 and Pregnant" on teen births, and exposure to Sesame Street on grade completion.

in 2000 on average (see Table 1), although low circulation counties had slightly lower employment, lower median income, higher poverty, and an older population. These level differences are not inherently problematic, as they will be controlled for by the origin/destination fixed effects. Of more concern to causal identification are changes over time that are correlated with pre-fracking circulation.

Columns (3) and (4) in Table 1 suggest that migration trends in low and high circulation counties are parallel between 2000 and 2010. Pre-2009 circulation rates do not predict changes in migration between 2000 and 2010, but do have predictive power for some local demographic trends. However, these differences are quite small: the predicted differences associated with an increase in readership from the 25th to the 75th percentile of *USA TODAY* circulation are never more than 0.51 percentage points. *New York Times* and *Wall Street Journal* pre-2009 circulation rates predict similarly small changes (Appendix Table A3).¹⁴ Even though these differences are small and unlikely to matter, I can also include origin county by year fixed effects to control for any origin-level characteristics like these that might be changing over time.

My baseline estimation corresponds to the thought experiment as follows

$$Y_{oSt} = \beta_1 newspaper\ exposure_{oSt} + \beta_2 newspaper\ exposure_{oSt}^2 + \phi_{oS} + \psi_{St} + \varepsilon_{oSt}. \quad (4)$$

The main outcome of interest is the inverse hyperbolic sine of the number of migrants from origin county o to fracking counties in state S in year t . The inverse hyperbolic sine approximates a natural log transformation but is defined for flows with zero migrants, allowing me to approximate the percent effect of news exposure. The state is used as the level of destination because few news articles reference specific counties by name, while state is frequently mentioned, meaning this is the level of variation. I only capture migration to fracking counties in the destination state, not the entire state, meaning I can examine flows from non-fracking counties to fracking counties in the same state. Origin counties with any fracking are excluded from the sample, as information in

¹⁴Readership of the *New York Times* and the *Wall Street Journal* are highly correlated, and the predicted effects are similar. The one characteristic that varies the most across newspapers is median household income. The *New York Times* and *Wall Street Journal* have higher readership in large urban areas that saw larger increases in earnings.

the news might affect the decisions of people originally living in fracking counties differently.¹⁵

Origin county by destination state fixed effects control for time-invariant pair specific characteristics, and destination state by year fixed effects control for changing characteristics of the fracking destinations. As suggested by the theoretical framework, I include news exposure quadratically to capture decreasing marginal returns to information, although the relationship is robust to news exposure entered linearly as well as different functional forms (see Appendix Table A6). To account for correlated shocks across geography, the standard errors are adjusted for clustering at the origin DMA level (203 clusters), a geographic measure meant to capture media markets. Observations are equally weighted.¹⁶

I begin with the specification in equation (4) because the identifying variation is highly transparent: origin counties experience different exposure to news about a specific destination because they have different pre-fracking circulation of national newspapers. I then progressively adjust this baseline specification to address potential concerns associated with this identifying variation. First, I include a vector of time varying origin county labor market controls, including the employment to population ratio, the unemployment rate, and average earnings (in 2010\$) to capture observable changes in the origin labor market. Second, I include origin county by year fixed effects which account for both observed and unobserved components of λ_{ot} . This is possible because I observe migration flows to 16 different fracking states from each origin county/year pair. Origin by year fixed effects control for changing characteristics of the origin county that affect migration flows. For example, if counties with higher circulation rates, and thus higher newspaper exposure, are changing over time in ways that reduce migration costs (e.g., becoming younger, more educated, or more wealth), these omitted variables might affect decisions to move to fracking areas in general. Origin county by year fixed effects absorb these and other changes over time and exploit variation in news exposure across potential destinations from the same origin. This makes a within origin county

¹⁵Including these counties does not significantly impact the results (see Appendix Table A7).

¹⁶If I instead weight by the origin county population in 2000, the impact from equation (4) is about 1.5 times as large and significant, but less precisely estimated. The loss in precision appears to be driven by the very largest counties. If I weight by the natural log of the origin county population in 2000 or weight by the origin county population in 2000, but exclude the top ten percent of counties by population, the estimate are similar in magnitude and precisely estimated.

comparison, to see if destination states that had more news exposure also experienced larger increases in migration flows. In this specification, any remaining confounding omitted variables must be origin/destination pair specific and vary over time (contained in ν_{ost}). For this reason I next include a similarly constructed measure of local news exposure.¹⁷ If local and national news exposure are strongly correlated and local news is endogenous to migration preferences, omitting local news will bias the coefficient on national newspaper exposure.¹⁸ I also conduct placebo tests and alternative estimation strategies designed to verify that the observed relationship is not driven by unobserved origin/destination specific changes over time.

As *newspaper exposure*_{ost} is a weighted sum, it is not immediate how to interpret the coefficients. If pre-2009 circulation rate in equation (3) is one (every household receives the newspaper) an additional news article will increase newspaper exposure by one unit. In reality, newspaper circulation rates are significantly lower than one hundred percent. I divide *newspaper exposure*_{ost} by 0.05, such that a one unit increase is equivalent to one additional news article in a newspaper with a five percent circulation rate. This level of circulation is comparable to a county with high readership of *USA TODAY*.¹⁹ Conveniently, when using this scaling average news exposure among treated observations is 0.99, suggesting a one unit increase also approximates the mean effect.

¹⁷This measure includes all domestic newspapers available through LexisNexis with available circulation data. Many local news sources provide free access to content online, which is not captured by this measure of local news exposure. National and regional news sources often provided limited free access, but ultimately require a paid subscription. The AAM circulation data includes digital replica newspapers, but not necessarily individual browsing behavior. To the extent that online exposure is positively correlated with print exposure, the estimates will simply represent the response to total news exposure (where print exposure is used as a proxy).

¹⁸The actual correlation between national and local newspaper exposure is 0.12.

¹⁹*USA TODAY* circulation ranges from 0 to 27.8 percent, with a mean of 1.2 percent; *New York Times* circulation ranges from 0 to 3.3 percent, with a mean of 0.51 percent; and the *Wall Street Journal* circulation ranges from 0 to 6.4 percent, with a mean of 1.2 percent. Even though average circulation is low, there is substantial variation, which is exploited by this identification strategy. The coefficient of variation is 0.90 for *USA TODAY*, 1.07 for the *New York Times*, and 0.52 for the *Wall Street Journal*.

V Results

V.A Graphical Analysis: Pre-trends and Treatment Effects by Circulation

First I present event study graphical evidence of the impact of national newspaper exposure on migration to verify that origin counties that will eventually be highly exposed to news do not have differential trends the news exposure, relative to origins that are less exposed. Consistent with my main strategy, I will focus on differences in initial circulation rates of national newspapers that will eventually report on fracking.²⁰ This tests to see if origin/destination specific news exposure is correlated with other unobserved characteristics that evolve over time and affect migration (ν_{oS_t}). For each origin county I collapse the pre-fracking circulation rates of the *USA TODAY*, *New York Times*, and *Wall Street Journal* to a single weighted average, where the weights are the share of the total national news articles about fracking in destination S in each newspaper. This measure captures the extent to which an origin will eventually be exposed to news about fracking in the destination state. I interact this measure with year indicators between 2001 and 2012 (omitting 2000 as the reference year), and regress the inverse hyperbolic sine of the number of migrants on this set of interactions to trace out the impact of pre-fracking circulation on migration over time as follows:

$$Y_{oS_t} = \sum_{\tau=2001}^{2012} \theta_{\tau} \text{Circulation}_{oS} * 1\{t = \tau\} + \phi_{oS} + \psi_{S_t} + \varepsilon_{oS_t}. \quad (5)$$

I include origin-destination pair fixed effects as well as destination-by-year fixed effects to exploit the same variation used in the main analysis. The coefficients on these year interactions are interpreted as the marginal effect of a one percentage point increase in the pre-fracking circulation rate on migration flows in that given year, and are plotted with 95 percent confidence intervals in Figure 3. For reference, a bar graph of the average number of articles about fracking in a specific state is superimposed, to show when news content about fracking was published.

Before 2008, only one 2006 *New York Times* article mentioned a specific destination state.

²⁰This measure is ideal for testing that different levels of initial circulation do not follow differential trends. The figure is almost identical when looking at alternative measures of treatment, such as the total newspaper exposure summed over all years.

Starting in 2008 there are small increases in the number of articles about fracking with a large jump in 2011 and 2012. Before 2010, migration fluctuates around zero, with only one statistically significant, negative estimate in 2003. Starting in 2006 there is a slight, insignificant upward trend, but overall it appears that origins that would eventually be highly exposed to news about fracking followed similar trends in migration. Since treatment also starts during this time, it is not clear this is evidence of non-parallel pre-trends.²¹ In 2010 the effect on migration becomes significant, and discontinuously jumps in 2011, when news content increased dramatically. The data suggest that a one percentage point increase in the pre-fracking circulation rate did not increase migration prior to news exposure, but was associated with a 2.5 percent increase in migration in 2011 and 2012, precisely when there was intense news coverage of fracking.²²

V.B Impact of Newspaper Exposure on Migration

I now report the regression results from equation (4) in Column (1) of Table 2. Given the absence of news in early years, I interpret effects as changes from zero to one. For an origin county with a five percent circulation rate, one additional newspaper article about fracking in a specific state increased migration flows to fracking counties in that state by 2.4 percent on average ($0.025 * 1 - 0.001 * 1$).²³ As average news exposure is also approximately one, this would suggest the mean effect of news exposure on migration was 2.4 percent as well. The most news was published in 2012 (average news exposure was 1.8) suggesting that in 2012, news about fracking increased migration flows to fracking counties by 4.2 percent on average.

I next adjust the baseline specification as outlined above to determine if changing characteristics of the origin bias the estimates. Controlling for annual origin county-level labor market measures in Column (2) does not change the coefficients. Including origin by year fixed effects in Column

²¹This pattern could also arise if origin destination pairs with high circulation followed an upward trend that was suppressed during the Great Recession (2007-2010), only to rebound in 2011. The pattern is essentially unchanged if I control for this by including origin county unemployment rates, employment to population ratios, and average earnings.

²²As seen in Appendix Figure A1, commuting responds similarly, although the increase is larger (8-12 percent) and begins earlier in 2009.

²³These estimates are not just statistically significant due to a large sample. As seen later, the significance remains when estimated over much smaller subsamples.

(3) absorbs the labor market measures included in Column (2) as well as any other unobserved characteristics of the origin that are changing over time and affect migration behavior. In this specification the effect of one additional newspaper article is 2.5 percent, and not statistically different from the baseline estimates. Finally, in Column (4) I include the origin by year fixed effects and control for local newspaper exposure. The effect of one additional national newspaper article remains 2.4 percent. For completeness, I repeat the same estimation using the number of migrants in levels as the outcome. In each of these specifications the marginal impact ranges from 1.4 to 1.7 and is not statistically distinguishable.²⁴ For the remainder of the paper, I will estimate the model corresponding to Column (2), which includes controls for labor market conditions at the origin, although the results are not sensitive to this choice of specification.

The data suggest that increased national news coverage of fracking increased migration to the fracking counties in states publicized. Although these estimated impacts are small, they are both statistically and economically significant. They imply that news about fracking increased migration flows to fracking counties by 2.4 percent on average. When considering the levels specification, exposure to news about fracking in a particular state led to 1.4-1.7 additional migrants from each origin on average. Another way to think about the effect size is to quantify how many people need to be exposed to the information in the news in order for one person to move. For reference the average county population in 2000 was 85,498 and the average household size was 2.59 people. If 5 percent of households receive the newspaper this would suggest that when 11,072 ($0.05 * 2.59 * 85,498$) people saw the article about fracking an additional 1.44-1.67 people would move, or only 0.015 percent. In other words, it takes about 10,000 interactions with the news about fracking for one additional person to move. In many cases, the news exposure measure is capturing more articles distributed to fewer readers. But even at a one percent circulation rate, we would still expect 2,214 ($0.01 * 2.59 * 85,498$) people to be exposed to multiple articles with 1.44-1.67 people moving, or only 0.075 percent. These ratios ignore any spillover effects of news passing through social networks, which would lead to even more people being exposed to the news. For a given origin this effect

²⁴An increase of 1.4 migrants represents a much larger effect at the mean than captured by the inverse hyperbolic sine specification. This appears to be driven by origin counties with large migrant flows. If the sample is restricted to origin/destination pairs with non-zero flows, the two specifications yield similar percent effects at the mean.

is small, but when aggregated up for a given destination the effect is large. This would suggest that providing information about potentially lucrative labor market opportunities elsewhere can increase migration to those destinations.

V.C Heterogeneous Responses Across Demographic Groups

The IRS data only provides the number of migrants and does not provide demographic characteristics. Exploring heterogeneous effects could shed light on who responds to the news and provide further evidence of the information mechanism. The annual American Community Survey (ACS) asks residents where they lived in the previous year, allowing me to construct origin/destination migration rates for demographic subgroups from the microdata. Unfortunately, the geographic data is only available starting in 2005 and the smallest geographic unit is the Public Use Micro Area (PUMA). PUMAs are geographic areas defined by population that are large enough to preserve privacy. Furthermore, migration geographic data is only available at the Migration PUMA (MIGPUMA) level, which are often even larger and contain one or more counties. There are several aspects of the data that are likely to make it more difficult to detect an effect. First, the ACS is a one percent sample of households so there is likely to be measurement error in the constructed migration rates which will reduce precision. Second, there is less geographic variation (and less variation in circulation rates) than is available at the county level leading to less precision.

Estimating an equation analogous to equation (4), I explore effect heterogeneity by gender, race, education, and marital status and plot the total effect of one unit of newspaper exposure for various subpopulations in Figure 4. When looking at the full population I estimate a significant 1.2 percent increase in migration to the state being mentioned in the news.²⁵ This effect is only half as large as the estimate using IRS data and less precise, as we would expect given the concerns about power discussed above. The impact for the full population is mostly driven by men, where I estimate a 0.95 percent increase in migration. Other groups that see significant increases in migration are 18-34 year olds, some college (at the ten percent level) and the unmarried. These are the same groups that largely drove the total migration response to fracking (Wilson, 2020). The impacts for high school graduates and dropouts (which also were drawn to fracking areas) are insignificant,

²⁵Estimates are similar if I include the local labor market controls.

but these groups are also less likely to be exposed to these news sources, which might explain the insignificant response.

Subgroups we would expect to be highly responsive, such as men with some college, unmarried men, young men (18-34), unmarried men with some college, young men with some college, and young unmarried men see significant effects ranging between 0.74 and 1.1 percent. This is highly consistent with both the data and anecdotal evidence about who moved to fracking booms. In other words, we see precisely the people we would expect to respond moving after the news.

V.D Impact of Newspaper Exposure on Cross-County Commuting

To avoid the monetary, psychic, and amenity costs that might accompany a move, an individual can choose to commute rather than migrate. Many people took advantage of the earnings gains associated with fracking by commuting rather than migrating (Wilson, 2020). Exposure to information might also affect this margin. In Table 3 I report the impact of newspaper exposure on the total number of workers who live in one county but work in a fracking county in the state mentioned in the newspaper article. For an origin county with a five percent circulation rate, one additional news article about fracking in a specific state increased the number of workers commuting to fracking counties in that state by approximately 6.6 percent. The impact on commuting is nearly three times as large as the migration response, which is not surprising as commuters avoid many of the fixed costs associated with moving.

When looking across the three pre-defined age groups, the response for one additional news article for 30 to 54 year olds is 5.2 percent and statistically larger than the response of both younger workers (3.1 percent) and older workers (3.6 percent).²⁶ This pattern is consistent with age-specific patterns in newspaper readership and commuting which offset each other. Newspaper readership (i.e., exposure) increases with age (Pew Research, 2013b), while geographic mobility falls with age (Molloy et al., 2011).²⁷ In Appendix Table A4, I also report differences by the

²⁶The reader will notice that the percentage effect is larger for all workers than for any of the three subgroups. This is in part because the pooled specification constrains the controls and fixed effects to be the same for each group. If this specification is run in levels, the effect for all workers is the sum of the effects for each subgroup, as expected.

²⁷In Figure 4 35-44 year olds exhibit a decline in migration. This is plausible as migration and

earnings and broad industry of the job they are commuting to. Consistent with people commuting to high paying fracking jobs, workers commuting to jobs that pay over \$3,333 a month (\$40,000 a year) are the most responsive. Consistent with previous work finding employment impacts across industries, I observe commute responses across industries, suggesting that news exposure not only induced people to commute for oil and gas extraction, but for other jobs affected by the labor market shock.

V.E Robustness

These estimates are robust to functional form and specification decisions. In Appendix Table A5, I re-estimate a variant of equation (7) including and excluding various fixed effects. The coefficients on migration are essentially unchanged by the inclusion of destination by year and origin by year fixed effects suggesting unobserved trends at the origin or destination do not play an important role. I also add origin state by destination state by year fixed effects, to compare migration flows from origin counties in the same state to a given fracking destination and evaluate the role of origin/destination pair specific characteristics. The coefficient on migration is essentially unchanged. The coefficients for commuting are more sensitive to the inclusion of various fixed effects, but are still significant even under the most conservative specifications. The effects of newspaper exposure on both migration and commuting are robust to including newspaper exposure linearly, quadratically, as a cubic, or as the inverse hyperbolic sine (see Appendix Table A6); excluding each destination state one at a time; sample and year restrictions (Appendix Tables A7 and A8);²⁸ or accounting for censoring in the IRS migration data (Appendix Table A9). I also explore differences across distance, by estimating equation (4) for origin county by destination state pairs in one hundred mile bins and plot the total marginal effect of news exposure for each distance in Appendix Figures A2 and A3. The effect climbs to about 6 percent between 400 and 1,000 miles and then gradually falls, consistent with information provision having no effect on people that might be aware of nearby opportunities, but a larger impact if it makes them aware of distant commuting are different decisions and potentially even substitutes.

²⁸Estimates are insensitive to including fracking origins, excluding non-fracking origins in fracking states, or limiting the sample to counties with positive newspaper circulation to focus on intensive margin variation in exposure.

potential opportunities.

The inverse hyperbolic sine specification is appropriate if information in the news has constant relative effects on the number of migrants. We might observe constant relative effects if, for example, treatment effects are heterogeneous and the information in the news interacts with other individual or local characteristics to effect migration. However, we might alternatively expect information to have constant absolute effects on migration rates. In Appendix Table A10, I show that news exposure significantly increases the number of migrants per person as well. One additional article in a newspaper with a 5 percent circulation rate leads to approximately 1 more migrant for each 200,000 people at the origin. These estimates imply similar exposure to migration ratios, suggesting that 10,000 interactions with the news induced 0.8-1.2 additional people to move. I also report estimates using the inverse hyperbolic sine of the migration rate. These estimates are more precise, suggesting the data is better fit by constant relative effects specifications. These estimates suggest that newspaper exposure increased migration rates by 1.9 percent at the mean and provide further evidence that information in the news increased migration.²⁹

One concern is that the readership of the *New York Times* and *Wall Street Journal* is on average more-educated, higher income, and older than the typical migrant to fracking areas (Wilson, 2020). However, readers of the *USA TODAY* look similar to typical fracking migrants. In 2007, 68 percent of *USA TODAY* readership was male, 66 percent were 25-54, 57 percent did not have a college degree, 31 percent had household incomes below \$50,000, and 26 percent were renters (see Appendix Table A11). In Table 4, I estimate the impact of exposure to news from each of these three newspapers separately. The estimated effects on migration and commuting are largest and most significant for news in the *USA TODAY*, with smaller effects from the *New York Times*, and very imprecise, insignificant effects from the *Wall Street Journal*. The same patterns hold by newspaper when looking at the ACS microdata (Appendix Figure A4). This lends further support to the information mechanism as the largest responses are driven by the news source that is more accessible to the typical migrant to fracking.³⁰

²⁹Heterogeneous effects by demographics are similar when using group-specific migration rates. Men, young adults, individuals with some college, the college educated, and pairwise combinations of these groups report significantly higher migration rates in response to news exposure.

³⁰In theory, one could scale up the migration effects by the demographic group specific probability

VI Alternative Strategy: Newspaper Market Border Comparison

Although origin level characteristics do not appear to introduce bias into the estimates in Table 2, it is possible that news exposure is correlated with other unobserved origin destination pair characteristics captured by ν_{ost} . For example, counties that have higher circulation of national news, and thus higher news exposure might, for unobserved reasons, also have a higher propensity to migrate to fracking areas when a boom hits. This could happen if, for example, origin counties more tied to the oil and gas industry also had higher readership of national newspapers, and thus higher exposure when these booms happened. To address this potential concern, I employ an alternative strategy that exploits variation in news exposure among neighboring counties. Using all domestic newspapers that had at least one article about fracking between 1999 and 2012 and had circulation data available, I construct geographic markets for each newspaper that capture the set of counties in the newspaper’s distribution network. For many local newspapers, distribution costs inhibit broad distribution and these markets are composed of a small group of adjacent counties around a central hub.³¹ I then identify counties on the border of this distribution network as well as contiguous counties that do not receive the newspaper and compare the effect of news articles, specific to that newspaper, on migration and commuting for counties on both side of the market border. This is done in a stacked regression as follows

$$Y_{ost} = \gamma_1 Articles_{nSt} * InMarket_{on} + \gamma_2 Articles_{nSt}^2 * InMarket_{on} + \gamma_3 InMarket_{on} + X'_{ot}\Gamma + \phi_{oS} + \psi_{nSt} + \varepsilon_{ost}. \quad (6)$$

The outcome in equation (6) is the same as in previous specifications. The variable *Articles* is the number of articles in newspaper n published in year t about fracking in state S (in units of ten), while *InMarket* is an indicator variable that equals one if the origin county o is in the market for newspaper n . To be specific n uniquely identifies each newspaper and the corresponding market border. So counties that do not receive newspaper n but are on the other side of the market

of reading the USA TODAY, like a first stage, to understand which groups are more responsive to the news. However, given potential information spillovers across groups it is not clear that this exclusion restriction would hold.

³¹Over 90 percent of these newspapers distribute to 40 counties or less.

border will also be assigned to border n as control counties. Time-varying origin controls and an origin-destination pair fixed effect are included. A newspaper-by-destination state-by-year fixed effect is also included, making this a comparison of flows to the same destination among counties along the same market border. Because I am making a pairwise cross-border comparison, I include every county pair along a newspaper market border, meaning a county-year observation may appear multiple times if it borders several counties across the market border.

The identifying assumption is that counties on either side of the newspaper’s market border would evolve similarly, but for the news coverage about fracking. Because counties are being compared to other local counties, similar preferences and propensities among these neighboring counties captured in ν_{ost} will be differenced out.³² This will identify the causal effect of news coverage as long as propensities to migrate to fracking during booms is similar within the county pair.

These results are reported in Table 5. Relative to no articles, ten news articles significantly increased migration by 5.6 percent in counties that received the newspaper, relative to their neighbors. There was also a similar significant effect on cross-county commuting (4.9 percent). In this sample, the average circulation rate among in-market border counties was slightly higher than the benchmark five percent at 5.5 percent, making it easy to compare the magnitude of this effect to the estimates from the previous specification. In the average in-market county with a newspaper circulation rate slightly higher than five percent, one additional local news article increases migration by approximately 0.6 percent.³³

³²As seen in Appendix Table A12, counties just inside these market boundaries and just outside the boundaries appear quite similar in 2000 and experienced similar trends on average between 2000 and 2010. Between 2000 and 2010, counties just outside the market see a slightly larger drops in employment, smaller increases in poverty, less growth in the minority population, and more population aging when controlling for newspaper border fixed effects. However, these differences are small, for example, the difference in the employment decline is only 0.2 percent of the mean.

³³The point estimates are similar if I exclude national newspapers and if I restrict the sample to only include one newspaper market border per county to ensure that counties only appear once. To do this I take the set of newspaper market borders each county belongs to, and restrict the sample to only include the newspaper market border that had the highest number of articles about fracking among these market borders.

VII Additional Explorations

VII.A News about Fracking in Another State

News about fracking could provide general information about the labor market impacts of fracking or specific information about where these labor market impacts are occurring. To evaluate the relative importance of these channels, I estimate how migration flows to fracking counties in a particular destination state respond to news about fracking in a different state. For example, observing that migration to North Dakota is less responsive to news about fracking in Pennsylvania than to news about fracking in North Dakota would suggest the location signal is important.

In practice, I randomly assign all observations indexed by S the fracking news exposure of S' , one of the other 15 fracking states. For example, all observations for the destination Arkansas might be randomly assigned the news exposure of North Dakota, while the observations of North Dakota might be randomly assigned the news exposure of Pennsylvania. I then estimate the regression similar to equation (4), but replace $News\ Exposure_{oSt}$ with the randomly assigned $News\ Exposure_{oS't}$, and calculate the marginal impact of a one unit increase in $News\ Exposure_{oS't}$ (i.e., one news article in a county with a five percent circulation rate). I repeat this 200 times and plot the histogram of potential impacts in Panel A of Figure 5, with the estimated effect using actual news exposure from Table 2 indicated. This is not a placebo test, as general information can plausibly affect the outcome.

The effect from Table 2 is larger than all but 5 of the repetitions (2.5 percent), suggesting the location signal is significant. The distribution of effects using randomly assigned news exposure is centered around 0.013, suggesting news about fracking in a different state has some positive predictive power. However, this 1.3 percent effect cannot be strictly attributed to general information about fracking. National newspapers report about multiple destinations, and news exposure is positively correlated across potential destination states. Among the 200 regressions, the average correlation coefficient between actual news exposure and randomly assigned news exposure was 0.44. To some degree, randomly assigned news exposure will proxy for actual news exposure, which might drive the estimated 1.3 percent effect.

For this reason I adjust the regression specification to include origin county by year fixed effects, and repeat the process 200 times. This specification absorbs average changes in news exposure at the origin county level and exploits destination specific deviations from the origin average. This specification looks to see if, for example, an origin that had unusually high exposure to news about fracking in North Dakota saw larger increases in migration to fracking counties in Arkansas. Randomly assigned news exposure no longer proxies for actual news exposure, but provides a test to determine if destination specific fluctuations in news exposure impact the corresponding migration flows. If the effect size from these regressions were comparable to the estimate in Table 2, we would be concerned the results are driven by things correlated with news exposure and pre-fracking circulation rates in general, not a causal effect of news content. This histogram of potential impacts is plotted in Panel B of Figure 5, along with the estimated effect from the origin county by year fixed effects specification from Table 2. The effects are centered around zero and are all smaller than the estimated effect in Table 2, suggesting that destination specific deviations in news exposure largely affect migration flows to the destination that is being mentioned. These results are consistent with the news about fracking conveyed information about where fracking is occurring; however, I cannot rule out that there are simply differences across destinations in the content of the news.³⁴

³⁴This is related to the advertising literature (Garthwaite, 2014), suggesting information could either lead to more migration overall (expansion) or shift people from other destinations (share stealing). Include total exposure to news about any of the 16 destination states has no effect on migration, but a small effect on commuting (see Appendix Table A13). Including the news exposure for the state that received the highest exposure within an origin and year has a positive effect, but does not reduce flows to other fracking destinations, suggesting news exposure led to expansion, rather than shifting. Information in the news also does not simply induce people to move to the closest fracking area. It is possible that migrants to fracking areas simply shifted from moving to other, non-fracking areas. However, when looking at the impact of total news exposure about fracking on migration to non-fracking areas, the coefficient is a small positive, suggesting the news does not shift migrants from non-fracking to fracking areas.

VII.B Migration Dynamics: Chain and Short-term Migration

Understanding the dynamic nature of the migration response and the incidence of chain migration or short-term migration can shed light on the size and welfare implications of these effects (Dustmann & Gorlach, 2016a; Dustmann & Gorlach, 2016b). Chain migration is prevalent in the international migration literature and news exposure could induce this behavior if migrants inform acquaintances at their origin about the situation. This might be particularly relevant if there are cross-person spillovers as people convey news to acquaintances. To test for chain migration, I regress migration flows on news exposure as well as lags of news exposure to see if news last year or two years ago impacts migration this year (see Appendix Table A14). An impact of lagged news exposure would capture both delayed moves and people who move after a friend or acquaintance is induced to move by the news (chain migration). When I regress migration on this year's news, last year's news, and news from two years ago, the impact of current news is not significantly different from the estimate in Table 2. The coefficient on last year's news is much smaller (0.7 percent) and marginally significant, with no significant impact of news from two years ago. This would suggest that current news has an impact on migration and there is some margin for a small, delayed moving response or chain migration. If this lagged response is due to chain migration, we would expect the effect to become insignificant if we control for last year's migration. However, when I control for last year's migration, the effects of news exposure are unchanged, suggesting the effect of lagged news exposure is likely coming from delayed migration rather than chain migration.³⁵

People that respond to the news might also move only temporarily before moving back. Because the IRS migration data only captures flows and does not follow individuals, I cannot directly examine the duration of migrants moves. However, by looking at reverse migration dynamically we can understand if short-term stints and return migration are common. From the IRS data I construct migration rates from fracking counties in state S to county o . These are the reverse

³⁵It should be noted that this lagged dependent model with origin/destination pair fixed effects is only consistent under strict assumptions. Using the ACS microdata I also estimate the effect of news exposure on the probability of joining an existing household of relatives or non-relatives to capture chain migration. There is no significant effect on joining an existing household and the estimates are close to zero (Appendix Figure A5).

flows used throughout the analysis. I then estimate how news exposure in origin county o about fracking state S affect migration from S to o in the future (see Appendix Table A15). Exposure to news about fracking does not lead to increased reverse migration in that same year. However, news exposure in the previous one or two years is associated with a significant increase in reverse migration. The coefficients are smaller than the direct effect on migration, but they are similar in magnitude suggesting many of the people who were induced to move to fracking areas by the news eventually move back. This is consistent with the prior work documenting both elevated inflows and outflows in fracking areas (Wilson, 2020).

VII.C Positive Labor Market versus Negative Environmental News

While some news references positive characteristics of fracking such as jobs, booms, or growth, other news discusses negative aspects such as pollution, health, dangers, and earthquakes. Individuals might respond differently to positive labor market news and negative environmental news. Given the importance of the location signal, it could be the case that even negative environmental news could impact people with pre-conceived beliefs about fracking but who were uncertain about where it was occurring. I parse each article for specific keywords such as “growth”, “boom”, “contaminat~”, and “earthquak~” to determine the positive and negative content of each article. These statistics are reported in Appendix Table A2. I then classify an article as positive if it has at least two positive mentions and has more positive mentions than negative. Negative articles are similarly defined. I then estimate the separate effect of positive and negative newspaper exposure on migration and cross-county commuting in Table 6.³⁶ Relative to no newspaper exposure in a county with a five percent circulation rate, one positive article significantly increased migration by 4.0 percent. The effect of negative news is still positive and significant, but less than half as large, suggesting even negative news provides information that induces migration.

In contrast to migration, positive labor market and negative environmental newspaper exposure affect commuting similarly. Unlike migrants, long distance commuters do not bear some of the costs associated with fracking such as potential home water contamination, earthquakes, or noise

³⁶Exposure to neutral articles with less than two positive and two negative keywords are not included in this regression. Specifications including all three levels are similar but less precise.

on residential streets. As commuters do not face many of these amenity costs, the location signal from negative environmental news might be just as effective as a location signal from positive labor market news.

VII.D The Role of Origin County Labor Market Conditions

In recent years there has been considerable concern about decreasing labor market fluidity and mobility, especially when it appears that people in weak labor markets could encounter more abundant opportunities elsewhere (Molloy, Smith, Trezzi, & Wozniak, 2016).³⁷ I next test if the impact of news exposure varies by labor market strength at the origin to understand if providing information is particularly impactful in weak economic areas. To do this, I estimate a variant of my main specification as follows

$$\begin{aligned}
Y_{oSt} = & \beta_1 \text{newspaper exposure}_{oSt} + \beta_2 \text{newspaper exposure}_{oSt}^2 \\
& + \beta_3 \text{newspaper exposure}_{oSt} * \text{emp/pop}_{ot-1} + \beta_4 \text{newspaper exposure}_{oSt} * \text{emp/pop}_{ot-1}^2 \\
& + \beta_5 \text{newspaper exposure}_{oSt}^2 * \text{emp/pop}_{ot-1} + \beta_6 \text{newspaper exposure}_{oSt}^2 * \text{emp/pop}_{ot-1}^2 \\
& + \beta_7 \text{emp/pop}_{ot-1} + \beta_8 \text{emp/pop}_{ot-1}^2 + X'_{ot} \Gamma + \phi_{oS} + \psi_{St} + \varepsilon_{oSt}.
\end{aligned} \tag{7}$$

Where emp/pop is the lagged county employment to population ratio, for the adult population. I then calculate the total effect of one unit of newspaper exposure, which is allowed to vary quadratically with the employment to population ratio and use the delta method to obtain standard errors (the corresponding coefficients are provided in Appendix Table A16). The effects on both migration and commuting are plotted in percentage points in Figure 6 for county employment to population ratios between 60 and 85 percent (approximately the 15th to 90th percentile). Both the migration and commute responses are larger for weaker economic areas. A one unit increase in newspaper exposure led to a 2.8 percent increase in migration from counties with a low employment to population ratio, but had a small, one percent impact on migration from counties with a high employment to population ratio. Low employment counties saw commute flows increase by nearly 8 percent for an additional news article, while counties with high employment saw increases closer to 2 per-

³⁷This topic has also come up in the popular press (Brooks, 2016; Cohen, 2016).

cent. Exposure to news about fracking in distant, potential labor markets had a larger impact on migration flows from economically weak areas, suggesting informational constraints might be a contributing factor to differences in migration behavior.³⁸

Using this specification, I estimate the implied impact of various news exposure treatments at different points in the origin employment to population ratio distribution as reported in Appendix Table A17. These simulations highlight the heterogeneous impacts of different potential policy interventions. Even increasing the exposure level by one unit (one additional article in a county with a five percent readership rate) has substantial impacts across the distribution. The impacts are the largest for counties with employment to population ratios below the mean, and monotonically decrease as the local labor market conditions improve. For reference, the average actual exposure level in 2012 was between 1.5 and 2.5 for all levels of labor market strength. Doubling or quadrupling exposure leads to large migration impacts for counties with weak labor markets, with a steep downward gradient. For origins near the mean employment to population level, migration flows would have been 6.4 percent lower if there had been no news about these local fracking booms.³⁹

The news exposure levels that maximize the migration impact are nearly 10 times as large, but result in migration responses that are twice as large in weak labor markets than in stronger labor markets. As differences in actual exposure are small it is likely that only a small part of the heterogeneous impacts can be explained by differential exposure. The information is most impactful for people living in counties with weak labor markets, where the expected gains to moving are largest.

This has several potential policy implications. Information provision policies could increase geographic mobility, potentially resulting in more beneficial labor market transitions (Molloy et al., 2016) and higher economic mobility (Chetty & Hendren, 2016). Providing a modest amount of information about potential labor market opportunities in other parts of the country to all counties will significantly increase migration to those regions. However, a government facing limited resources would see the largest returns by focusing on providing information to weak labor markets. Not only would the migrant benefit, by encountering more favorable labor markets, but this might

³⁸The pattern is more flat, but still downward sloping if instead the unemployment rate is used.

³⁹The patterns are similar when looking at commute behavior although the impacts are larger and more heterogeneous, while the maximizing level are more uniform (see Appendix Table A18).

also generate positive externalities for workers in the weak origin labor market, as the market becomes less slack.

VII.E Impact of TV News Exposure

Over the years adults rely less on newspapers as their main source of news, while the television remains an important source for 69-74 percent of adults, and the internet has become increasingly important (Pew Research, 2013b). Data constraints prevent me from comparing internet news exposure to traditional news sources, but I am able to compare migration and commute responses to television and newspaper news exposure.

Using abbreviated news transcripts from the Vanderbilt Television News Archive (VTNA) for the three major TV news networks (ABC, CBS, and NBC) I construct a measure of TV news exposure similar to the measure of national newspaper news

$$TV\ news\ exposure_{ost} = \sum_{c \in C} (\text{broadcasts on } c \text{ about fracking in } S)_t * pre09\ view\ rate_{oc}. \quad (8)$$

The set $C = \{ABC, CBS, NBC\}$ and captures TV news coverage from the major national news networks. As with $newspaper\ exposure_{ost}$, $TV\ news\ exposure_{ost}$ captures variation in national news, which is weighted by the channel's pre-fracking Nielsen's viewership rates obtained through the 2008 Television and Cable Factbook. These pre-fracking viewership rates might be problematic. Between 2007 and 2009, TV stations transitioned from analog to digitally transmitted broadcasts on a market-by-market basis. When a market transitioned, viewers were required to obtain digital reception equipment. This might have induced some viewers to substitute to other outlets (i.e., cable), meaning viewership rates in 2008 might be less correlated with viewership rates at the time of the broadcasts for markets that transitioned after 2008. This might introduce measurement error in the viewership rates as only some markets had transitioned by the time data was collected. For this reason, I also run specifications using ratings from the latest 2016 Factbook, after all markets were updated.⁴⁰ Nielsen ratings are only available at the DMA-level, which is a mutually exclusive

⁴⁰Using 2016 ratings potentially introduces endogeneity if viewership is responding to migration and commute behavior. However, circulation rates are highly persistent, suggesting this bias might be small.

set of contiguous counties that represent a media market. As such, I aggregate up migration flows, newspaper exposure, and labor market measures from the county to the DMA-level.⁴¹ Typical viewership of ABC, CBS, and NBC was approximately 50 percent during this time period, so I scale TV news exposure such that a one unit increase represents the effect of one additional news broadcast from a TV network with 50 percent viewership.⁴²

DMA-level estimates are presented in Table 7. I first report the effects of newspaper exposure on migration as the level of analysis has changed. One additional news article in a DMA with a five percent circulation rate increased migration to the fracking state mentioned by 5.0 percent. This point estimate is twice as large as the county-level estimate, but is not statistically different. Column (2) reports the estimated effects for TV news exposure. Using 2008 viewership rates, TV news exposure had no significant impact on migration behavior. In Column (3) I regress migration on both newspaper exposure and TV news exposure.⁴³ The coefficients for both news sources remain similar and the effect for newspaper exposure is significant. In Columns (4) and (5) I conduct the same analysis, but use the 2016 viewership rates to measure TV news exposure in an attempt to reduce measurement error. When estimating the impact of TV news exposure alone one additional broadcast in a market with 50 percent viewership leads to a 8.7 percent (0.111-0.024) increase in migration. When we include both newspaper and TV news the effect of TV news is no longer significant but of a similar magnitude (the p-value on the first order effect is 0.11), suggesting TV news exposure might also affect migration, but the relationship is weaker than for newspaper news. When looking at commute behavior, both newspapers and TV news have separate, significant effects.⁴⁴

⁴¹There are only 203 DMA, which reduces the variation in the explanatory variable because circulation rates are now calculated over larger areas.

⁴²This measure does not capture cable news channels, such as CNN or Fox News. The VTNA only collects one hour of news broadcast data from these channels, and cable circulation is measured differently than traditional TV. If *TV news exposure_{ost}* is negatively correlated with cable news exposure (i.e., if network and cable news are substitutes) and both sources of news lead to more migration, than these estimates will be biased downward. If instead network and cable news are complements, network news could be interpreted as a proxy for total TV news.

⁴³Newspaper exposure and TV news exposure are moderately, positively correlated ($\rho = 0.36$).

⁴⁴When aggregating to the DMA-level, many neighboring counties fall into “fracking” DMAs that are excluded from the sample of origin DMAs. This likely attenuates the estimated impact on

The effects of TV news are larger in magnitude, but not significantly different. There are several potential explanations why. As explained above, the TV data might introduce more measurement error. In addition to data-related issues, there are also aspects of the TV news treatment that might lead to different effects. In general there are far fewer TV news broadcasts about fracking (only 17 total) distributed to a larger group of people. If the accumulation of information signals is a driving mechanism (which the conceptual framework would suggest) then a small group of people getting multiple pieces of information might have a larger effect than many people getting fewer pieces of information. This is consistent with a significant impact of newspaper exposure but a much weaker impact of TV news exposure. If the decision to commute is less costly than the decision to move, we might expect limited information (such as that provided in TV broadcasts) to be more likely to affect commuting, which is once again consistent with what we observe. Also, the nature of the information appears to be different across newspapers and TV news. The 17 TV news broadcasts are all short, only 1-5 minutes and provide less information than a newspaper article. If the quality of the information signal is lower, we might also be less likely to observe a response. Observing similar impacts for both newspaper and TV exposure would suggest that both the intensity and penetration of content influence the magnitude of the effect.

VIII Online Activity, a Potential Mechanism

The data indicate that exposure to news about fracking in a certain state, increases migration and commuting to fracking areas in that state. This relationship posits that news coverage provides information about potential labor market opportunities, and affects migration through changing expectations and uncertainty. This channel cannot be directly tested in the data. However, I can quantify how interest in fracking and the states mentioned changes after news is disseminated. Using Google Trends data, I explore search interest before and after TV news broadcasts about fracking.⁴⁵ For a specified search term (i.e., “fracking”), Google Trends will provide a time-series of search intensity at the national, state, or DMA level. This time-series is an ordinal measure of commuting.

⁴⁵Ideally I would also like to look at search behavior after newspaper articles are published. However, as there are over 560 articles, the pre- and post- windows for each article overlap extensively.

intensity that equals 100 on the day with the highest number of searches per capita, and with every other day scaled as a percent of the maximum. For example, on a day that is assigned a value of 20, search intensity for the search term was only 20 percent the level from the maximum day. This measure facilitates comparisons within a region over time, but is not conducive to studying differences across both geography and time. As such, I will examine changes in search behavior a short period before and after a TV news broadcast, but cannot reliably determine if search intensity increased by more in areas with higher TV viewership rates.

For each of the 17 TV news broadcasts that mention “fracking” or “shale”, I pull daily time-series for every DMA in the United States for 15 days before the broadcast, the day of the broadcast, and 14 days after for several search terms.⁴⁶ First I look at search intensity for the term “fracking”, and then I look at search intensity for the name of any states that are mentioned in the broadcast. States are only mentioned in 14 of the broadcasts. To identify the impact of the news broadcast on average search intensity I estimate the following regression

$$search\ index_{opt} = \sum_{\tau=-14}^{14} \delta_{\tau} * \mathbf{1}\{t\text{ is } \tau \text{ days from broadcast}\}_{op} + X_t'\Gamma + \phi_{op} + DOW_t + \varepsilon_{opt} \quad (9)$$

where $search\ index_{opt}$ is the search index on date t in DMA o relative to the search period p . The search period is the 15 days prior, the day of, and the 14 days after each broadcast, such that op uniquely identifies each DMA/period pair, over which the relative search index is measured. The set of coefficients δ_{τ} trace out the daily deviations in the search index from the omitted day ($\tau = -15$). I include a DMA by search period fixed effect in order to compare days from the same search that have comparable indices. Day of the week fixed effects are also included to account for differences in search behavior during different times in the week.⁴⁷

⁴⁶A special thanks to Tanner Eastmond for help working through the Python code.

⁴⁷Several of the broadcasts were in close proximity to other high publicity events connected to either fracking or the states mentioned in the reports. For example, On January 24, 2012, four days prior to a news report about fracking in Pennsylvania, President Barack Obama discussed shale gas extraction and fracking in the State of the Union. Similarly, late on December 31, 2011, there was an earthquake in Youngstown, Ohio, that many linked to fracking, just four days before a news report on fracking in Ohio. Other high publicity state specific events (such as college football bowl games or school shootings) also occur during some of the search period windows. When looking at

These effects are plotted in Figure 7 for “fracking” and for the state names. Search intensity for “fracking” spikes the day of the broadcast and remains elevated for the next two days before falling back to the previous levels. Across all DMAs, search intensity for “fracking” jumps by nearly 2.5 points on average. Because the search index is a relative measure, this cannot be converted to back out how many additional searches were made. If I combine days into 3 day bins, for statistical power, I estimate a 2.5 point spike in days 0 to 2, followed by an statistically significant one point increase for the remaining 12 days in the sample, suggesting that search interest remained elevated for some time (see Appendix Figure A7).

When looking at search interest in the names of states that were mentioned in the broadcast there is also a spike one day after the broadcast and interest remains elevated for the next five days. For reference I also regress the search intensity for the same set of fracking states that are not mentioned in the news broadcasts. These estimates remain close to zero, with no spike or increase after the news broadcast, suggesting this is not capturing overall interest in fracking states. Although not direct evidence that news coverage motivates people to move, the Google Trends analysis suggests news coverage induces people to seek more information about the potential fracking destination.

Using Twitter content, I can also examine if the news impacts people’s tweet behavior. To do this I pull all tweets that reference a specific term (e.g., “fracking”) within the 30-day broadcast window, I then plot the average daily number of tweets across all of the broadcast event.⁴⁸ As with Google Trends, there is a spike in tweets, re-tweets, and likes of tweets that include “fracking” in the days following a broadcast (see Figure 8). In the appendix I show that these tweet often include links or mention specific Twitter users by their handle. There are also spikes in positive terms like “jobs”, “money”, “opportunity”, and “pay” (see Appendix Figures A8-A10).

searches for fracking I include indicator controls for the three days after the two events related to fracking, and when looking at searches for specific state names I include a set of indicator controls for the local, high publicity events that are listed in Appendix Table A19. If I do not control for these events the series becomes more volatile, but there is still a significant spike directly after the broadcast (Appendix Figure A6).

⁴⁸Because I am accessing historic data I cannot use the Twitter API, so I can only access location for a limited number of users. As such, this analysis only includes one series for each broadcast event and examines the interrupted time-series.

IX Conclusion

Migration is a way for individuals to improve the types of labor markets they encounter, but many of those that appear to face the largest benefit do not move. In this paper I evaluate the role of information in the decision to move to labor market opportunities. The current literature speaks very little to the effect of labor market information on migration behavior.

To estimate the effect of news on migration I exploit information about local fracking booms disseminated through the national news. I combine national news content with historic local circulation rates to construct a measure of news exposure that strips away endogenous changes in consumer readership and endogenous changes in producer content decisions.

The data suggest that for a county with a five percent circulation rate and no previous exposure, one news article about fracking in a specific state increased migration flows to fracking counties in that state by 2.4 percent. Cross-county commute flows also increase, by 6.6 percent. Fracking news about a specific destination increases migration to that state, but not other fracking states, suggesting the location signal is important. There is evidence that news exposure induces short-term migration, with many people returning after several years. Migration flows are more responsive to exposure to positive labor market news than negative environmental news, though both lead to more migration. In contrast, commute flows respond similarly to positive and negative news, consistent with commuters not facing many of the negative costs associated with fracking at their homes (e.g., water contamination, increased risk of earthquakes). Other sources of news, such as TV news exposure, has an effect on commuting and potentially migration also. As further evidence that news coverage increases interest in these fracking destinations, I find that, directly after a TV news broadcast about fracking, Google search interest in both the term “fracking” and the names of states mentioned significantly increases and tweets from Twitter also respond. News exposure induces people to seek more information on the internet, which might influence the migration decision.

Importantly, the migration response is largest from origin counties that have been experiencing weak labor market conditions, suggesting the benefit to news provision is largest in those areas. This has potential implications when trying to understand why less-educated and low-income households

in poor performing labor markets are unlikely to move, and if there are policies that can encourage more migration to better economic opportunity. All else equal, providing more information about potential labor market opportunities in other areas would increase geographic mobility in all areas, with the most pronounced response in weak labor markets where the returns to migration are the largest.

References

- Bartik, Alexander, Janet Currie, Michael Greenstone, and Christopher Knittel. 2017. "The Local Economic and Welfare Consequences of Hydraulic Fracturing," Available at SSRN: <https://ssrn.com/abstract=2692197>.
- Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 5(2): 9-49.
- Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F. Katz, and Christopher Palmer. 2019. "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," NBER Working Paper No. 26164, National Bureau of Economic Research, Cambridge, MA.
- Bottan, Nicolas and Ricardo Perez-Truglia. 2017. "Choosing Your Pond: Revealed-Preference Estimates of Relative Income Concerns," Working Paper No. 23615, National Bureau of Economic Research, Cambridge, MA.
- Bound, John and Harry Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s," *Journal of Labor Economics*, 18(1): 20-54.
- Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: the Case of Seasonal Migration in Bangladesh," *Econometrica*, 82(5): 1671-1748.
- Brooks, Arthur C. 2016. "How to Get Americans Moving Again," *The New York Times*, May 20, 2016. <https://www.nytimes.com/2016/05/21/opinion/how-to-get-americans-moving-again.html>.
- Bureau of Labor Statistics, U.S. Department of Labor. 2015. The Economics Daily, State employment-population ratios in 2014 on the Internet at <https://www.bls.gov/opub/ted/2015/state-employment-population-ratios-in-2014.html> (visited July 18, 2017).
- Cascio, Elizabeth and Ayushi Narayan. 2020. "Who Needs a Fracking Education? The Educational Response to Low-Skill Biased Technological Change," *ILR Review* accepted.
- Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz. 2015. "The Effect of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 106 (4): 855-902.
- Chetty, Raj, Nathaniel Hendren. 2016. "The Impact of Neighborhoods on Intergenerational Mobility II: County-Level Estimates," Working Paper No. 23002, National Bureau of Economic Research, Cambridge, MA.
- Chong, Alberto and Eliana La Ferrara. 2009, "Television and Divorce: Evidence from Brazilian "Novelas"," *Journal of the European Economic Association*, 7(2/3): 458-468.
- Cohen, Patricia. 2016. "Fewer Americans Strike Out for New Jobs, Crimping the Recovery," *The New York Times*, May 24, 2016. <https://www.nytimes.com/2016/05/21/opinion/>

- how-to-get-americans-moving-again.html.
- DellaVigna, Stefano and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting," *Quarterly Journal of Economics*, 122(3): 1187-1234.
- Dionne, Georges and Scott E. Harrington. 1991. An Introduction to Insurance Economics, in Foundations of Insurance Economics: *Readings in Economics and Finance*, Georges Dionne and Scott Harrington eds. (Boston, Mass: Kluwer Academic, 1991).
- Dustmann, Christian and Joseph-Simon Gorlach. 2016. "The Economics of Temporary Migrations," *Journal of Economic Literature*, 54(1): 98-136.
- Dustmann, Christian and Joseph-Simon Gorlach. 2016. "Estimating Immigrant Earnings Profiles when Migrations are Temporary," *Labour Economics*, 41: 1-8.
- Farre, Lidia and Francesco Fasani. 2013. "Media exposure and internal migration – Evidence from Indonesia," *Journal of Development Economics*, 102: 48-61.
- Feyrer, James, Erin Mansur, and Bruce Sacerdote. 2017. "Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution," *American Economic Review*, 107(4): 1313-1334.
- Ganong, Peter and Daniel Shoag. 2017. "Why has Regional Income Convergence in the U.S. Declined?" Working Paper No. 23609, National Bureau of Economic Research, Cambridge, MA.
- Garthwaite, Craig and Timothy J. Moore. 2012. "Can Celebrity Endorsements Affect Political Outcomes? Evidence from the 2008 US Democratic Presidential Primary," *The Journal of Law, Economics, & Organization*, 29(2): 355-384.
- Garthwaite, Craig. 2014. "Demand Spillovers, Combative Advertising, and Celebrity Endorsements," *American Economic Journal: Applied Economics*, 6(2): 76-104.
- Gentzkow, Matthew. 2006. "Television and Voter Turnout," *Quarterly Journal of Economics*, 121(3): 931-972.
- Gentzkow, Matthew, and Jesse Shapiro. 2010. "What Drives Media Slant? Evidence from U.S. Daily Newspapers," *Econometrica*, 70(1): 35-71.
- Greenwood, Michael J. 1975. "Research on Internal Migration in the United States: A Survey," *Journal of Economic Literature*, 13(2): 397-433.
- Greenwood, Michael J. 1985. "Human Migration: Theory, Models, and Empirical Studies," *Journal of Regional Science*, 25(4): 521-544.
- Hanson, Gordon and Craig McIntosh. 2010. "The Great Mexican emigration," *the Review of Economics and Statistics*, 92(4): 798-810.
- Jensen, Robert and Emily Oster. 2009. "The Power of TV: Cable Television and Women's Status in India," *Quarterly Journal of Economics*, 124(3): 1057-1094.
- Kaplan, Greg and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration," *International Economic Review*, 58(1): 57-94.
- Kearney, Melissa and Phillip B. Levine. 2015a. "Early Childhood Education by MOOC: Lessons from Sesame Street," Working Paper No. 21229, National Bureau of Economic Research, Cambridge, MA.
- Kearney, Melissa and Phillip B. Levine. 2015b. "Impact of MTV's *16 and Pregnant* on Teen Childbearing," *American Economic Review*, 105(12): 3597-3632.
- Kearney, Melissa and Riley Wilson. 2018. "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom," *Review of Economics and Statistics*, 100(4): 678-690.
- Kling, Jeffrey, Jeffrey B. Liebman, and Lawrence F. Katz. 2007. "Experimental Analysis of Neighborhood Effects," *Econometrica*, 75(1): 83-119.

- La Ferrara, Eliana, Alberto Chong, and Suzanne Duryea. 2012. "Soap Operas and Fertility: Evidence from Brazil," *American Economic Journal: Applied Economics*, 4(4): 1-31.
- Langley, Paul C. 1974. "The Spatial Allocation of Migrants in England and Wales: 1961-66," *Scottish Journal of Political Economy*, 21: 259-277.
- Malamud, Ofer and Abigail Wozniak. 2012. "The Impact of College Migration: Evidence from the Vietnam Generation," *Journal of Human Resources*, 47(4): 913-950.
- Maniloff, Peter and Ralph Mastromonaco. 2014. "The Local Economic Impacts of Hydraulic Fracturing and Determinants of Dutch Disease," Division of Economic and Business Working Paper Series, Colorado School of Mines.
- McCauley, Jeremy. 2019. "The Role of Information in Explaining the Lack of Welfare-Induced Migration," *Working Paper*, University of Bristol.
- McKenzie, David, and Hillel Rapoport. 2007. "Network Effects and the Dynamics of Migration and Inequality: Theory and Evidence From Mexico," *Journal of Development Economics*, 84(1): 1-24.
- McKenzie, David and Hillel Rapoport. 2010. "Self-Selection patterns in Mexico-U.S. migration: the role of migration networks," *the Review of Economics and Statistics*, 92(4): 811-821.
- Molloy, Raven, Christopher Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States," *Journal of Economic Perspectives*, 25(2): 1-42.
- Molloy, Raven, Christopher Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding declining fluidity in the U.S. Labor Market," *Brookings Paper in Economic Activity*, Spring 2016.
- Monras, Joan. 2015. "Economic Shocks and Internal Migration," *IZA DP No.* 8840.
- Munshi, Kaivan. 2003. "Networks in the modern economy: Mexican migrants in the US labor market," *Quarterly Journal of Economics*, 118, 549-599.
- Notowidigdo, Matthew. 2013. "The Incidence of Local Labor Demand Shocks," Working Paper No. 17167, National Bureau of Economic Research, Cambridge, MA.
- O'Connell Paul G. J. 1997. "Migration Under Uncertainty: 'Try Your Luck' or 'Wait and See'," *Journal of Regional Science*, 37(2): 331-347.
- Pew Research Center, Washington D.C. 2013a. "As Gas Prices Pinch, Support for Oil and Gas Production Grows, Those Aware of Fracking Favor Its Use," (March, 19, 2012). <http://www.people-press.org/2012/03/19/as-gas-prices-pinch-support-for-oil-and-gas-production-grows/>.
- Pew Research Center, Washington D.C. 2013b. "Amid Criticism, Support for Media's 'Watchdog' Role Stands Out," (August, 8, 2013). <http://www.people-press.org/2013/08/08/amid-criticism-support-for-medias-watchdog-role-stands-out/2-main-source-for-news/>.
- Pierce Kevin. 2015. "SOI Migration Data: A New Approach, Methodological Improvements for SOI's United States Population Migration Data, Calendar Years 2011-2012," The Internal Revenue Service.
- Rothschild, Michael, and Joseph Stiglitz. 1970. "Increasing risk: I. A definition," *Journal of Economic Theory*, 2(3): 225-243.
- Ruggles, Steve, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2015. Integrated Public Use Microdata Series: Version 6.0. Minneapolis: University of Minnesota.
- Shultz, Theodore W. 1961. "Investment in Human Capital," *American Economic Review*, 51 (1):1-17.
- Sjaastad, Larry. 1962. "The costs and returns to human migration," *Journal of Political Economy*, 70(supplemental): 80-93.
- Tobin, James. 1965. "Money and Economic Growth," *Econometrica*, 33(4): 671-684.

- Todaro, Michael P. 1969. "A Model of Labor Migration and Urban Unemployment in Less Developed Countries," *American Economic Review*, 59(1): 138-148.
- Wilson, Riley. 2020. "Moving to Economic Opportunity: the Migration Response to the Fracking Boom," *the Journal of Human Resources*, accepted.
- Winters, Paul, Alain de Janvry, and Elsabeth Sadoulet. 2001. "Family and Community Networks in Mexico-U.S. Migration," *the Journal of Human Resources*, 36(1): 159-184.
- Wiswall, Matthew and Basit Safar. 2015. "How Do College Students Respond to Public Information about Earnings," *Journal of Human Capital*, 9(2): 117-169.
- Wozniak, Abigail. 2010. "Are College Graduates More Responsive to Distant Labor Market Opportunities?" *Journal of Human Resources*, 45(4): 944-970.

Tables

Table 1

County Characteristics by *USA TODAY* Pre-Fracking Circulation Rate

	County Characteristics in 2000		Change from 2000 to 2010		Predicted Difference from 25th to 75th Percentile
	Below Median (1)	Above Median (2)	Below Median (3)	Above Median (4)	
<i>Migrants to fracking areas (Percent of Population)</i>	0.09	0.11	-0.01	-0.02	0.00
<i>Employment to Population (16+)</i>	55.68	59.14	-0.78	-1.43	-0.10
<i>Unemployment Rate</i>	3.29	3.44	0.96	1.33	-0.09
<i>Median Household Income</i>	31,805	38,834	8,485	8,763	290
<i>Percent in Poverty</i>	15.82	12.29	0.83	2.18	0.51***
<i>Percent White</i>	85.16	84.32	-0.83	-2.26	-0.58***
<i>Percent Black</i>	8.95	9.27	-0.15	0.42	0.17***
<i>Percent Hispanic</i>	5.87	5.72	1.77	2.38	0.29***
<i>Percent Other Race</i>	5.89	6.41	0.98	1.85	0.41***
<i>Percent Population 20-34</i>	16.67	19.60	-0.55	-0.67	0.02
<i>Percent Population 35-64</i>	38.76	38.56	1.82	1.47	-0.27***
<i>Percent Population Over 64</i>	15.96	13.64	1.20	1.15	0.00
<i>Percent Households Renting</i>	23.47	28.30	1.64	1.81	0.10
<i>Number of Counties</i>	1,420	1,418	1,420	1,418	2,838

Notes: Migration data from the IRS Statistics of Income. Other county characteristics obtained through American FactFinder from the 2000 Census and 2010 Census and 5-Year American Community Survey. *USA TODAY* circulation data from the Alliance for Audited Media. The county level median pre-2009 circulation rate of the *USA TODAY* was 0.83 percent, and ranged from 0 to 27.8 percent. Median Household Income is reported in current dollars. Column (5) reports the predicted change in the characteristic between 2000 to 2010 when pre-2009 circulation increases from the 25th to the 75th percentile. Standard errors are corrected for clustering at the origin DMA level. p<0.01 ***, p<0.05 **, p<0.1 *.

Table 2

Impact of Destination State Specific National Newspaper Exposure on Migration to Fracking Counties in State

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}				Number of Migrants _{oSt}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>National Newspaper Exposure_{oSt}</i>	0.025*** (0.004)	0.025*** (0.004)	0.026*** (0.004)	0.025*** (0.004)	1.439*** (0.419)	1.491*** (0.437)	1.673*** (0.534)	1.644*** (0.532)
<i>National Newspaper Exposure_{oSt}²</i>	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.036* (0.020)	-0.038* (0.020)	-0.036* (0.022)	-0.038* (0.021)
<i>Local Newspaper Exposure_{oSt}</i>				0.009** (0.004)				-0.671 (1.662)
<i>Local Newspaper Exposure_{oSt}²</i>				-0.0001** (0.00004)				0.057 (0.043)
<i>Origin Labor Market Controls</i>		X				X		
<i>Origin by Year Effects</i>			X	X			X	X
<i>Origin/Destination Local News</i>				X				X
<i>Mean Number of Migrants</i>	7.6	7.6	7.6	7.6	7.6	7.6	7.6	7.6
<i>Observations</i>	590,224	590,224	590,224	590,224	590,224	590,224	590,224	590,224

Notes: Data from the IRS Statistics of Income, LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For origin/destination pairs with any news exposure, mean national newspaper exposure is 0.99. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Origin controls include the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$). Origin/destination specific local news is all destination state specific fracking news content listed in LexisNexis from non-national domestic newspapers. The variable *Local Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 40 percent circulation rate, approximately the 95th percentile of pre-fracking circulation among non-national newspapers with articles about fracking. The sample correlation between national news exposure and local news exposure is approximately 0.12. Origin county by year fixed effects control for time-varying characteristics of the origin county and account for potential changes in preferences toward fracking that might be correlated with newspaper readership and affect migration to fracking areas. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table 3

Impact of Destination State Specific Newspaper Exposure on Cross-County Commuting to Fracking Counties in State

	Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt} By Age			
	All Jobs (1)	Under 30 (2)	30-54 (3)	Over 54 (4)
<i>National Newspaper Exposure_{oSt}</i>	0.068*** (0.009)	0.032*** (0.005)	0.053*** (0.007)	0.037*** (0.005)
<i>National Newspaper Exposure_{oSt}²</i>	-0.002*** (0.0004)	-0.001*** (0.0002)	-0.001*** (0.0003)	-0.001*** (0.0002)
<i>Dependent Mean</i>	31.4	8.6	18.0	4.9
<i>Observations</i>	499,440	499,440	499,440	499,440

Notes: Data from the LEHD Origin-Destination Employment Statistics (LODES), LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2002 to 2012. LODES data is only available starting in 2002. Observations are not limited to origin counties in the same local commuting zone and includes long distance commuters. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For comparison, circulation of the *USA TODAY* was 4.5 percent at the 95th percentile. Commuting jobs are also examined by age groups, pre-defined in the LODES data. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. The effect of national newspaper exposure on commuting is significantly larger for workers aged 30-54 than the other two age groups. p<0.01 ***, p<0.05 **, p<0.1 *.

Table 4

Heterogeneity by Newspaper: Impact of Newspaper Exposure on Migration and Commuting to Fracking Regions

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}		Inverse Hyperbolic Sine of the Number of Cross-County Commuting Jobs _{oSt}	
	(1)	(2)	(3)	(4)
<i>USA TODAY Exposure_{oSt}</i>	0.027*** (0.004)	0.021*** (0.004)	0.081*** (0.013)	0.030*** (0.007)
<i>USA TODAY Exposure_{oSt}²</i>	-0.001*** (0.0003)	-0.001*** (0.0002)	-0.003** (0.001)	-0.001 (0.001)
<i>New York Times Exposure_{oSt}</i>	0.007*** (0.002)	0.009*** (0.002)	0.019*** (0.004)	0.016*** (0.004)
<i>New York Times Exposure_{oSt}²</i>	-0.0001** (0.0001)	-0.0002*** (0.0001)	-0.001*** (0.0002)	-0.0004** (0.0002)
<i>Wall Street Journal Exposure_{oSt}</i>	0.098* (0.055)	0.056 (0.046)	0.097 (0.130)	0.018 (0.086)
<i>Wall Street Journal Exposure_{oSt}²</i>	-0.027 (0.036)	-0.006 (0.029)	-0.062 (0.098)	-0.033 (0.060)
<i>Origin by Year Effects</i>		X		X
<i>Dependent Mean</i>	7.6	7.6	31.4	31.4
<i>Observations</i>	590,224	590,224	499,440	499,440

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis Newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012 for the migration data and 2002 to 2012 for the commute data. Each newspaper's exposure level is scaled to represent the impact of one additional news story in a county with circulation at the 95th percentile (3.9 percent for the *USA TODAY*, 1.9 percent for the *New York Times*, and 2.4 percent for the *Wall Street Journal*). Controls include the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$). In all specifications origin/destination pair fixed effects and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table 5

Newspaper Market Cross Border Analysis: Impact of Newspaper Articles on Migration and Commuting

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}			Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>News articles_{nSt} * In-Market_{on}</i> <i>(in 10s of Articles)</i>	0.061*** (0.013)	0.056*** (0.012)	0.069*** (0.017)	0.054** (0.023)	0.042** (0.021)	0.082*** (0.032)
<i>News articles_{nSt}² * In-Market_{on}</i> <i>(in 10s of Articles)</i>	-0.005*** (0.001)	-0.004*** (0.001)	-0.005*** (0.002)	-0.005** (0.002)	-0.004** (0.002)	-0.007** (0.003)
<i>In-Market_{on}</i>	-0.001*** (0.0002)	-0.001*** (0.0002)	-0.00004** (0.00002)	-0.001** (0.0004)	-0.001* (0.0005)	-0.0001* (0.00003)
<i>Exclude National Newspapers</i>		X			X	
<i>Only One Border per Origin County</i>			X			X
<i>Dependent Mean</i>	24.4	24.0	14.1	144.0	141.8	72.2
<i>Observations</i>	1,476,352	1,465,648	509,664	1,250,112	1,240,784	431,360

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012 for the migration data and 2002 to 2012 for the commute data. Sample includes all counties on both sides of the border of a newspaper market for any of the 220 newspapers with an article about fracking and circulation data. I include every county pair along a newspaper market border, meaning a county-year observation may appear multiple times if it borders several counties across the market border. News articles are newspaper and destination state specific and measured in units of ten. In-market is an indicator that equals one if the county is inside the newspaper's market area (i.e., has positive circulation). In all specifications origin/destination pair fixed effects are included to control for time-invariant differences across pairs. Newspaper market border by destination by year fixed effects are also included to control for characteristics of the local border/destination that vary over time, and make this a comparison of origin counties within the same newspaper market border. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. Average circulation among in-market counties across all newspapers was 5.5 percent. Columns (2) and (5) exclude national newspapers, as their market borders are not local. In Columns (3) and (6) the sample is restricted to only include one newspaper market border per county, and it is the border that had the highest number of articles about fracking among all of the borders the county belongs to. Standard errors adjusted for clustering at the origin designated market area are in parentheses. p<0.01 ***, p<0.05 **, p<0.1 *.

Table 6

Positive Labor Market vs. Negative Environmental News: Impact of Newspaper Exposure on Migration and Commuting

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}		Inverse Hyperbolic Sine of the Number of Cross-County Commuting Jobs _{oSt}	
	(1)	(2)	(3)	(4)
<i>Positive Newspaper Exposure_{oSt}</i>	0.044*** (0.010)	0.047*** (0.011)	0.072*** (0.018)	0.056*** (0.020)
<i>Positive Newspaper Exposure_{oSt}²</i>	-0.004** (0.002)	-0.004** (0.002)	-0.007** (0.003)	-0.003 (0.003)
<i>Negative Newspaper Exposure_{oSt}</i>	0.026*** (0.005)	0.025*** (0.005)	0.098*** (0.014)	0.051*** (0.011)
<i>Negative Newspaper Exposure_{oSt}²</i>	-0.001*** (0.0005)	-0.001*** (0.0004)	-0.006*** (0.001)	-0.002*** (0.001)
<i>Origin by Year Fixed Effects</i>		X		X
<i>Dependent Mean</i>	7.6	7.6	31.4	31.4
<i>Observations</i>	590,224	590,224	499,440	499,440

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012 for the migration data and 2002 to 2012 for the commute data. *Exposure_{oSt}* measures are scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. A positive news article is one that contains at least two positive phrases (referencing jobs, boom, or growth) and more positive than negative phrases (referencing pollution, health, danger, or earthquakes), while a negative article is the opposite. Some fracking destinations have many positive and negative articles, leading to a high correlation between *Positive Newspaper Exposure_{oSt}* and *Negative Newspaper Exposure_{oSt}* ($\rho = 0.70$). Controls include the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$). In all specifications origin/destination pair fixed effects and destination by year fixed effects, are included to control for time-invariant differences across pairs and characteristics of the destination and origin that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *.

Table 7

Source of News: Impact of Newspaper and TV News Exposure on Migration to Fracking Regions

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}					Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Newspaper Exposure_{oSt}</i>	0.052*** (0.018)		0.052*** (0.018)		0.051** (0.018)	0.045** (0.019)		0.055** (0.019)		0.045** (0.019)
<i>Newspaper Exposure_{oSt}²</i>	-0.002*** (0.001)		-0.002*** (0.001)		-0.002*** (0.001)	-0.001** (0.001)		-0.001** (0.001)		-0.001** (0.001)
<i>TV News Exposure_{oSt}</i>		0.040 (0.072)	0.029 (0.072)	0.111* (0.065)	0.100 (0.064)		0.125** (0.062)	0.115* (0.061)	0.144** (0.069)	0.134* (0.069)
<i>TV News Exposure_{oSt}²</i>		-0.010 (0.017)	-0.011 (0.017)	-0.024* (0.014)	-0.021 (0.013)		-0.019 (0.014)	-0.019 (0.014)	-0.021 (0.014)	-0.018 (0.014)
<i>2008 TV Viewership Rates</i>		X	X				X	X		
<i>2016 TV Viewership Rates</i>				X	X				X	X
<i>Dependent Mean (in Levels)</i>	60.2	60.2	60.2	60.2	60.2	152.6	152.6	152.6	152.6	152.6
<i>Observations</i>	32,864	32,864	32,864	32,864	32,864	27,808	27,808	27,808	27,808	27,808

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. TV news circulation is only available at the Designated Market Area (DMA) level, from the 2008 Television Factbook, and all data is aggregated to that level. The level of observation is the origin DMA by destination state by year from 2000 to 2012. The variable *Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For comparison, circulation of the *USA TODAY* was 4.5 percent at the 95th percentile. The variable *TV News Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional TV news broadcast on a network with a 50 percent circulation rate, approximately the average circulation rate of ABC, CBS, or NBC. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Controls for the origin DMA unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In 2008, there was significant transition to digital TV and full viewership ratings were not available, so Columns (4),(5), (9), and (10) use TV circulation from 2016 to construct TV news exposure. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Figures




Figure 1

Fracking Counties and Shale Plays

Notes: Any county with production from fracking wells between 2000 and 2012 is labeled as a fracking county. Shale play boundaries are outlined in black.

Source: Author's calculations constructed from DrillingInfo well level data. Shale play boundaries are from the EIA.

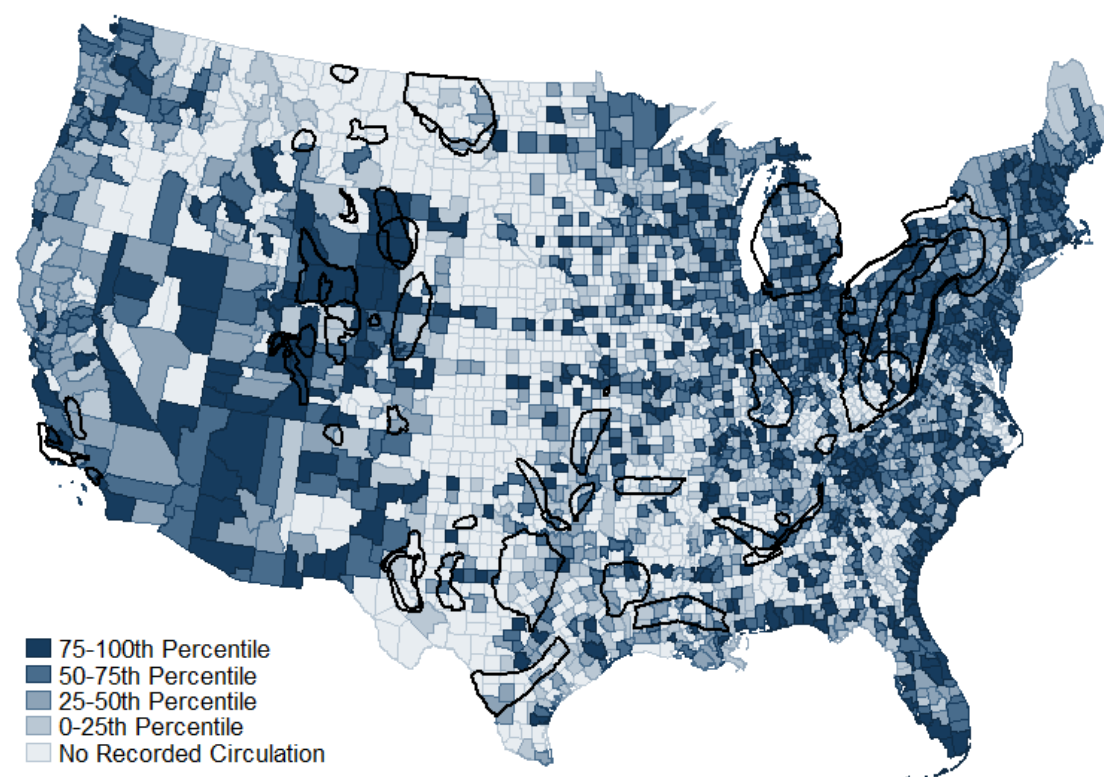


Figure 2

County-level Circulation of USA TODAY between 2005 and 2008

Notes: Location of shale plays outlined in black.

Source: Author's calculations using annual county-level circulation rates averaged between 2005 and 2008 obtained from the Alliance of Audited Media.




Figure 3

Trends in Migration by Pre-fracking Circulation

Notes: For each origin, the pre-fracking circulation rate is the weighted average of the pre-fracking circulation of the *USA TODAY*, *New York Times*, and *Wall Street Journal*, where weights are the share of the total articles about fracking in the destination state in each newspaper. This measure captures the extent to which an origin will eventually be exposed to fracking news. This measure is then interacted with year indicators. The year 2000 is treated as the base year. The inverse hyperbolic sine of the number of migrants is then regressed on this set of interactions along with origin-destination pair effects and destination-by-year fixed effects, as in the main specification, to trace out the effect of a one percentage point increase in the pre-fracking circulation rate on migration, as a percent. The coefficients on these year interactions are interpreted as the marginal effect of a one percentage point increase in the pre-fracking circulation rate on migration flows in that given year and are plotted for each year on the right axis, to look at trends by differences in eventual exposure. Standard errors are corrected for clustering at the origin DMA level, with 95 percent confidence intervals provided. For reference, the average number of articles about fracking in each state is also plotted for each year in bars on the left axis.

Source: Author's calculations using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and migration flows from the IRS SOI.

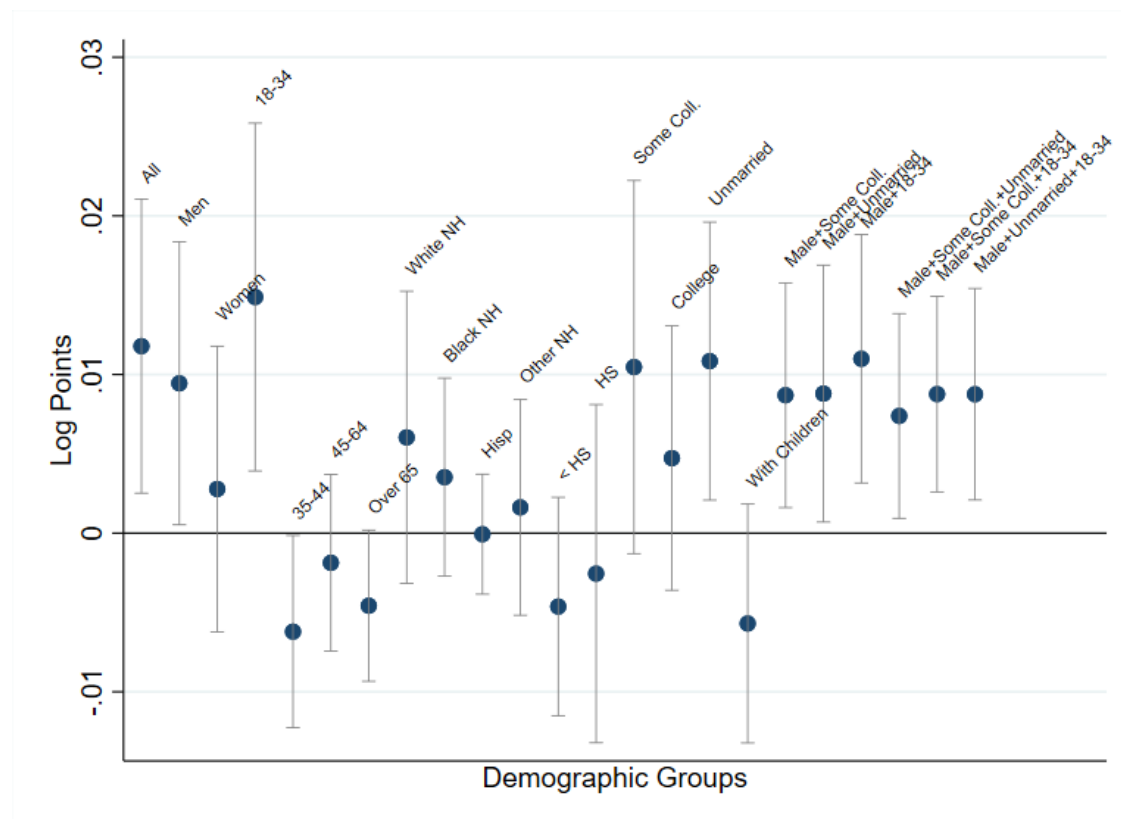


Figure 4

Migration Impact of Newspaper Exposure by Demographic Group

Notes: The total impact of a one unit increase in Newspaper Exposure is plotted with 95 percent confidence intervals. The level of observation is the origin MIGPUMA by destination state by year from 2005 to 2012 for the given group. Origin MIGPUMAs with any fracking production are excluded. The variable *National Newspaper Exposure_{ost}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. Origin/destination pair fixed effects and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the origin state.

Source: Author's calculation using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and microdata from the 2005-2012 American Community Survey.

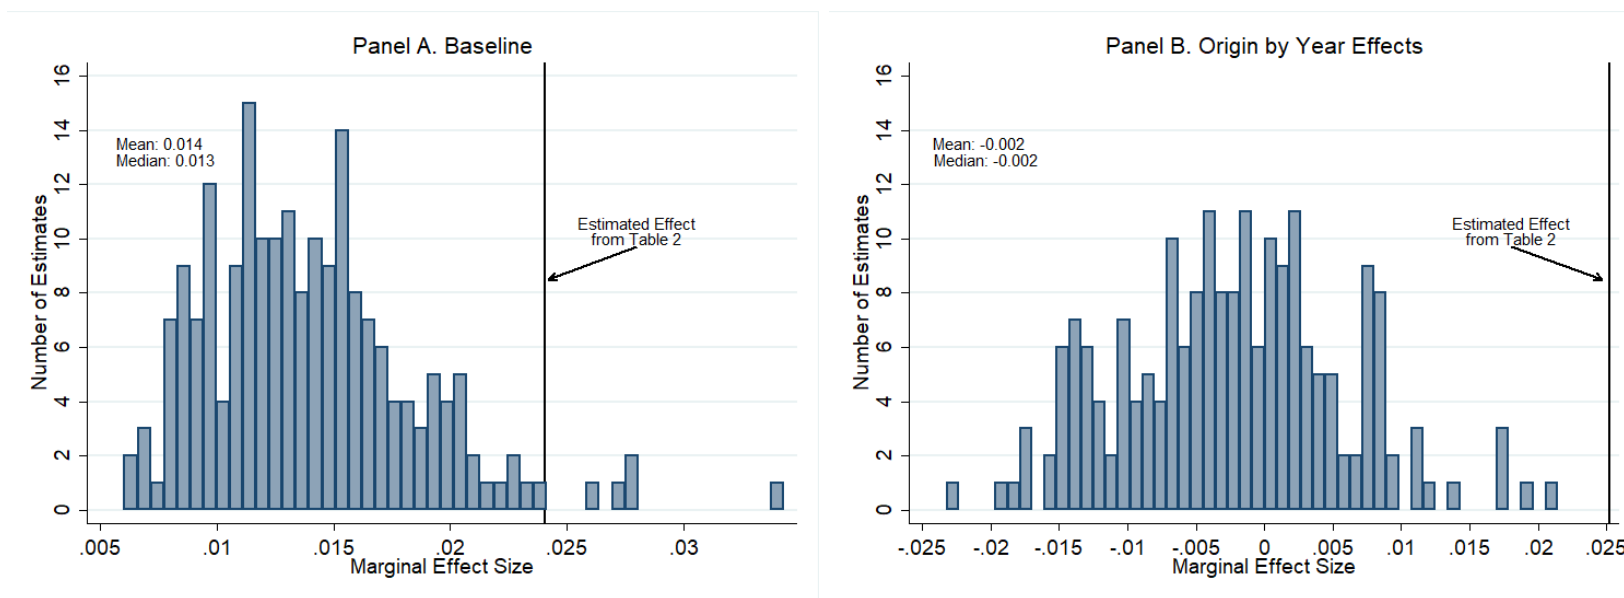


Figure 5

Location Signals: Migration Response to Randomly Assigned News about Fracking in a Different State

Notes: Each state is randomly assigned the fracking news exposure of a different state, and then the inverse hyperbolic sine of migration is regressed on a quadratic of this randomly assigned news exposure, similar to the baseline regression in equation (4). The histogram of estimated effects from the baseline model for 200 regressions are plotted in Panel A. For some states the trends in news exposure are similar, and across all 200 regressions the average correlation between actual news exposure and randomly assigned news coverage was 0.44. Panel B. repeats the same 200 regressions but includes origin by year effects. This exploits variation in news coverage across destinations within an origin, relying on destination state specific deviations in news exposure.

Source: Author's calculation from 200 regressions of randomly assigned news exposure on the inverse hyperbolic sine of migration using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and county to county migration flows from the IRS SOI.

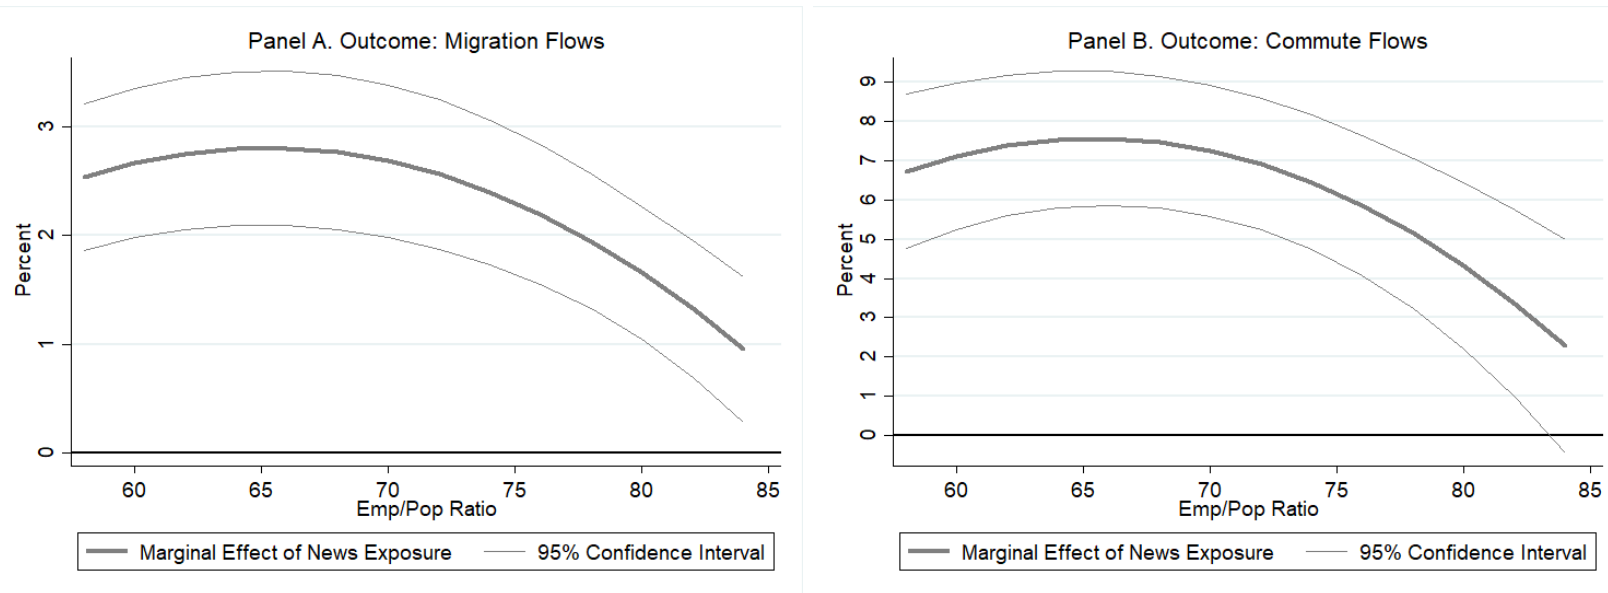


Figure 6

Heterogeneous Impacts of Newspaper Exposure by Origin Employment to Population Ratio in $t - 1$

Notes: Marginal impact of newspaper exposure calculated by interacting a quadratic in newspaper exposure and a quadratic of lagged employment to population ratio at the origin. Approximately the 10th to 90th percentile of the employment to population ratio are plotted. Standard errors are calculated using the delta method.

Source: Author's calculations using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, migration flows from the IRS SOI, and county employment to population ratio constructed from BLS QCEW data.

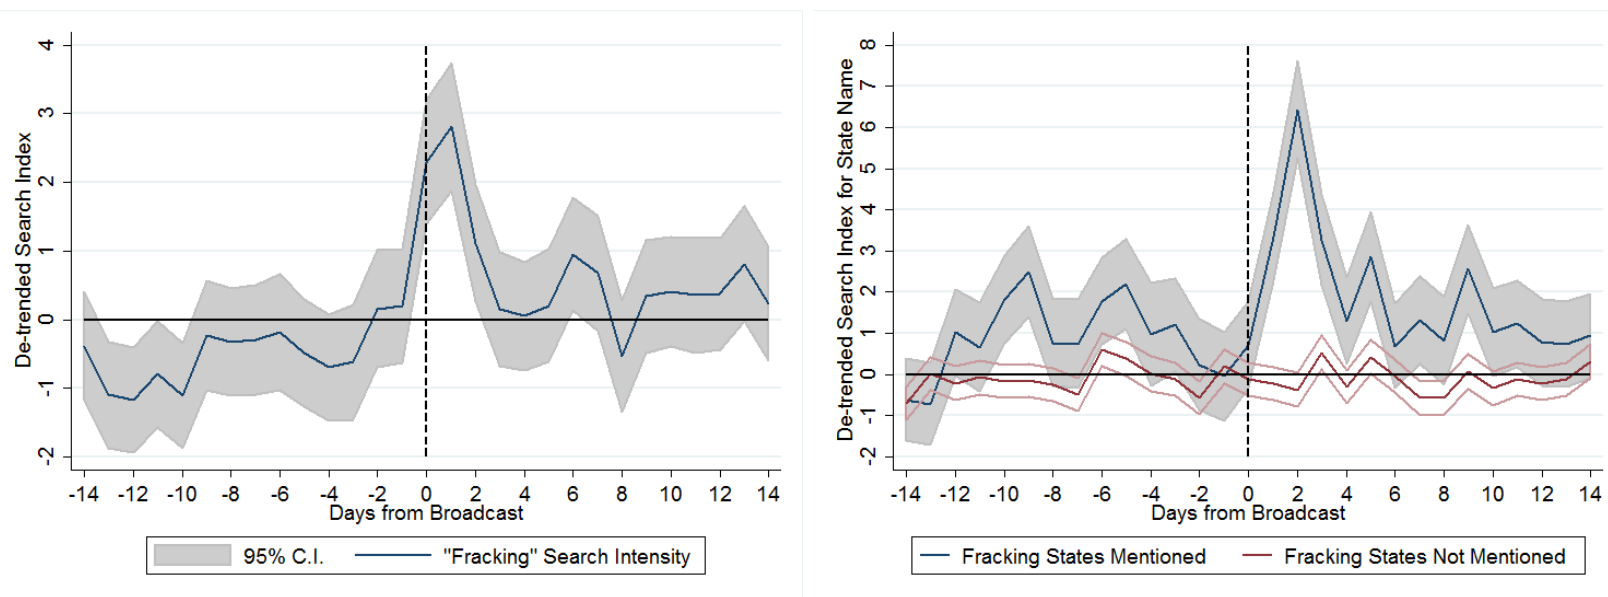


Figure 7

Google Search Interest in “Fracking” and Names of Fracking States Mentioned in TV News Broadcasts

Notes: Plot depicts the average daily search index for the term “fracking” by DMA before and after 17 TV broadcast mentioning fracking or shale gas between 2006 and 2012 as recorded by the Vanderbilt Television News Archive. Search intensity is de-trended by removing day of week and search (DMA by four week publication window) specific effects. To be consistent with other analysis in the paper, one broadcast from CNN and one broadcast from Fox News are excluded. Four days prior to a news broadcast on January 28, 2012, President Barack Obama mentioned shale gas exploration due to fracking in the State of the Union Address. Four days prior to a news broadcast on January 4, 2012, there was an earthquake in Ohio that reporters linked to fracking. For both of these event I include indicator variables for the next four days. Additional control indicators are also included for specific high publicity state-specific events that fall in the search period window, such as the earthquakes, wildfires, special elections, and major sporting events. Excluding these controls does not significantly change the daily average search index time series (see Figure A6). For reference, the search intensity for fracking states *not* mentioned in the news broadcast is also plotted with 95 percent confidence intervals. Standard errors are clustered at the search level.

Source: Source: Author’s calculations using daily search indices from Google Trends.

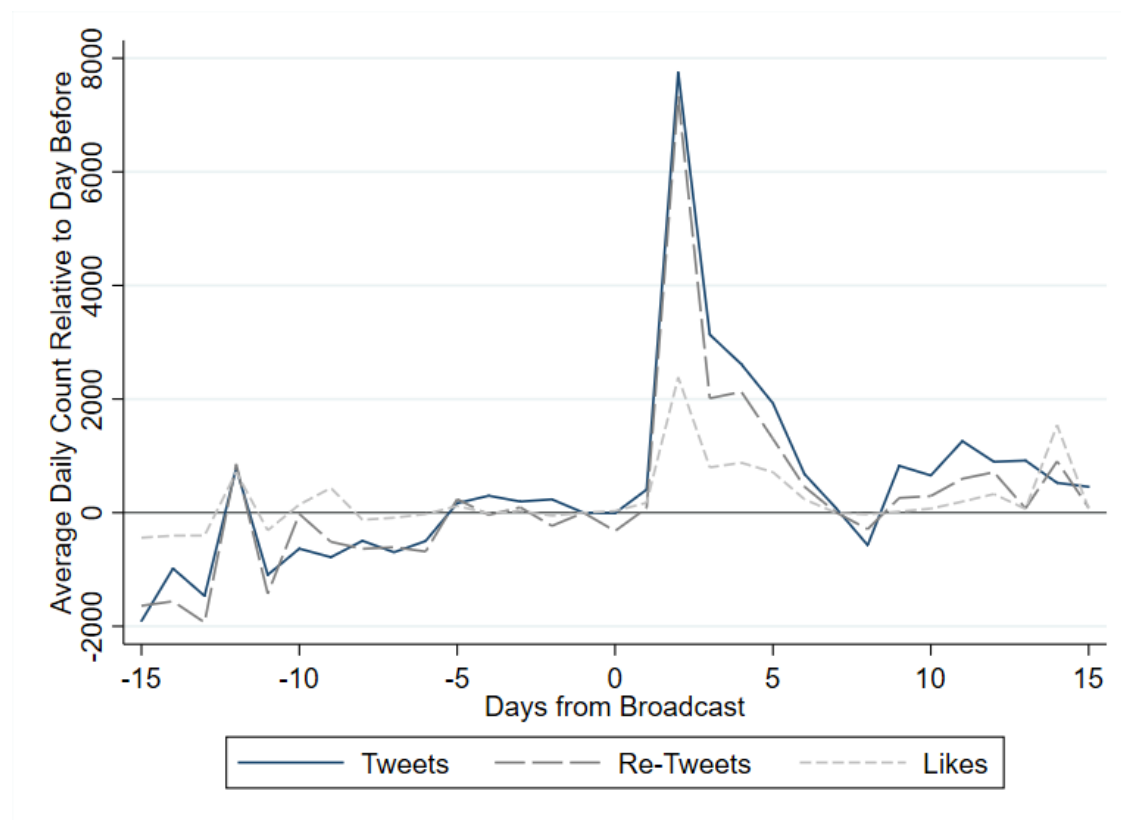


Figure 8

Average Tweets, Re-tweets, and Likes of Tweets about Fracking around a TV News Broadcast about Fracking

Notes: All tweets with the word “fracking” within 15 days of the post-2009 TV news broadcast events are included. The daily number of total tweets, re-tweets, and likes are averaged over all of the broadcast events. The Twitter-user’s location is not attainable, and I am unable to exploit geographic variation in exposure to news about fracking.

Source: Author’s calculations using TV news content from VTNA and Twitter tweet content.

Table A1

How Exposure to News about Fracking Impacts Awareness and Approval of Fracking

	All (1)	Male (2)	Female (3)	Under Median Age (4)	Over Median Age (5)	HS or Less (6)	Some College (7)	College Degree (8)
Outcome: Heard About Fracking								
<i>Newspaper Exposure In 2011 (State-level)</i>	0.01*** (0.003)	0.01* (0.005)	0.01*** (0.003)	-0.00 (0.007)	0.01*** (0.004)	0.01** (0.005)	-0.00 (0.006)	0.01** (0.005)
<i>Constant</i>	0.62*** (0.017)	0.65*** (0.032)	0.62*** (0.021)	0.56*** (0.022)	0.69*** (0.025)	0.49*** (0.026)	0.64*** (0.037)	0.79*** (0.020)
<i>Observations</i>	1,492	502	939	582	892	465	437	586
Outcome: Heard A lot About Fracking								
<i>Newspaper Exposure In 2011 (State-level)</i>	0.02*** (0.003)	0.02*** (0.004)	0.02*** (0.004)	0.02*** (0.005)	0.02*** (0.004)	0.02*** (0.005)	0.01** (0.006)	0.03*** (0.006)
<i>Constant</i>	0.24*** (0.021)	0.25*** (0.033)	0.25*** (0.023)	0.20*** (0.023)	0.28*** (0.026)	0.11*** (0.017)	0.23*** (0.033)	0.41*** (0.033)
<i>Observations</i>	1,492	502	939	582	892	465	437	586
Outcome: Approve of Fracking								
<i>Newspaper Exposure In 2011 (State-level)</i>	0.01*** (0.003)	0.03*** (0.006)	0.00 (0.003)	0.02** (0.007)	0.00 (0.004)	0.00 (0.006)	0.02*** (0.005)	0.01 (0.010)
<i>Constant</i>	0.51*** (0.022)	0.47*** (0.035)	0.52*** (0.023)	0.49*** (0.038)	0.53*** (0.032)	0.56*** (0.042)	0.53*** (0.034)	0.44*** (0.034)
<i>Observations</i>	1,030	357	652	346	673	252	300	477

Notes: Observations at the individual level from the Pew Research Center March 2012 Political Survey. Individuals were asked if they had heard “a lot”, “a little”, or “nothing at all” about fracking. Individuals who reported hearing “a lot” or “a little” about fracking were then asked if they favor or oppose fracking. As such I am unable to estimate how hearing about fracking affects people’s approval of fracking. I can only estimate the impact of exposure to the news on approval, conditional on hearing about fracking. Newspaper exposure constructed from LexisNexis newspaper content and AAM newspaper circulation. The median age in the sample is 48. Observations are weighted using the weights provided by the Pew Research Center. Standard errors are corrected for clustering at the state-level. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A2
Content of Newspaper Articles

	Share of Articles that Mention							
	<i>Jobs</i>			<i>Pollution</i>				<i>Total</i>
	<i>References</i> ¹	<i>“boom”</i>	<i>“growth”</i>	<i>References</i> ²	<i>“health”</i>	<i>“danger”</i>	<i>“earthquake”</i>	<i>Articles</i>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>All Articles</i>	0.17	0.24	0.14	0.39	0.20	0.12	0.06	562
<i>Mention State:</i>								
<i>Arkansas</i>	0.29	0.50	0.0	0.64	0.43	0.07	0.36	14
<i>California</i>	0.28	0.41	0.34	0.41	0.31	0.13	0.03	32
<i>Colorado</i>	0.10	0.39	0.23	0.65	0.28	0.13	0.06	31
<i>Louisiana</i>	0.32	0.47	0.26	0.58	0.21	0.26	0.05	19
<i>Michigan</i>	0.57	0.43	0.43	0.29	0.0	0.14	0.14	7
<i>Mississippi</i>	0.0	0.50	0.0	0.75	0.75	0.50	0.0	4
<i>Montana</i>	0.67	0.67	0.67	0.17	0.17	0.0	0.0	6
<i>New Mexico</i>	0.0	1.0	0.0	0.0	0.0	0.0	0.0	2
<i>North Dakota</i>	0.29	0.57	0.25	0.46	0.29	0.11	0.04	28
<i>Ohio</i>	0.33	0.36	0.12	0.55	0.39	0.21	0.24	33
<i>Oklahoma</i>	0.28	0.60	0.36	0.72	0.32	0.28	0.12	25
<i>Pennsylvania</i>	0.30	0.41	0.14	0.66	0.30	0.23	0.07	91
<i>Texas</i>	0.22	0.42	0.27	0.41	0.24	0.11	0.09	105
<i>Utah</i>	0.20	0.40	0.60	0.40	0.60	0.20	0.20	5
<i>West Virginia</i>	0.25	0.46	0.13	0.63	0.54	0.13	0.04	24
<i>Wyoming</i>	0.04	0.36	0.14	0.68	0.25	0.21	0.04	28

Notes: Newspaper content for articles between 2008 and 2012 obtained through LexisNexis, for the *New York Times*, *USA TODAY*, and *Wall Street Journal*. Not all articles reference a state, and some articles reference multiple states. Search terms are truncated to include various tenses and included both capitalized and lower case. ¹ Jobs References include the following search terms: “new job”, “creat~ + job”, “low + unemploy~”, “hire/hiring”. ² Pollution References include the following search terms: “contaminat~” and “pollut~”.

Table A3

County Characteristics by the *New York Times* Pre-Fracking Circulation Rate

	Pre-2009 Circulation Rate of the <i>New York Times</i>				
	County Characteristics in 2000		Change from 2000 to 2010		Predicted Difference from 25th to 75th Percentile
	Below Median (1)	Above Median (2)	Below Median (3)	Above Median (4)	
<i>Migrants to fracking areas (Pct. of Population)</i>	0.11	0.08	-0.02	-0.01	0.00
<i>Employment to Population (16+)</i>	55.62	59.22	-0.6	-1.62	-0.15
<i>Unemployment Rate</i>	3.50	3.23	0.82	1.47	0.17***
<i>Median Household Income</i>	31,652	39,018	8,080	9,173	1,119***
<i>Percent in Poverty</i>	16.05	12.04	1.11	1.90	0.07
<i>Percent White</i>	84.49	85.0	-1.26	-1.84	-0.41***
<i>Percent Black</i>	9.40	8.82	0.11	0.16	0.10***
<i>Percent Hispanic</i>	5.35	6.24	1.83	2.32	0.28**
<i>Percent Other Race</i>	6.11	6.19	1.15	1.68	0.31***
<i>Percent Population 20-34</i>	17.55	18.72	-0.40	-0.82	-0.13***
<i>Percent Population 35-64</i>	38.28	39.04	1.65	1.64	0.03
<i>Percent Population Over 64</i>	15.65	13.95	1.08	1.28	0.07
<i>Percent Households Renting</i>	25.37	26.40	1.75	1.69	-0.21***
<i>Number of Counties</i>	1,426	1,412	1,426	1,412	2,838

Notes: Migration data from the IRS Statistics of Income. Other county characteristics obtained through American FactFinder from the 2000 Census and 2010 Census and 5-Year American Community Survey. Circulation data for the *New York Times* from the Alliance for Audited Media. The county level median pre-2009 circulation rate of the *New York Times* was 0.32 percent, ranging from 0 to 3.29 percent. Circulation of the *New York Times* and the *Wall Street Journal* are highly correlated ($\rho = 0.8$), and characteristics look similar by circulation of the *Wall Street Journal*. Median Household Income is reported in current dollars. Column (5) reports the predicted change in the characteristic between 2000 to 2010 when pre-2009 circulation increases from the 25th to the 75th percentile. Standard errors are corrected for clustering at the origin DMA level. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *.

Table A4

Impact of Destination State Specific Newspaper Exposure on Cross-County Commuting to Fracking Regions

	Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}					
	By Monthly Earnings			By Broad Industry		
	≤\$1,250 (1)	\$1,250–\$3,333 (2)	≥\$3,333 (3)	Goods Producing (4)	Trade and Transportation (5)	Other Industry (6)
<i>National Newspaper Exposure_{oSt}</i>	0.034*** (0.005)	0.045*** (0.006)	0.045*** (0.006)	0.025*** (0.005)	0.029*** (0.004)	0.058*** (0.007)
<i>National Newspaper Exposure_{oSt}²</i>	-0.001*** (0.0002)	-0.001*** (0.0003)	-0.001*** (0.0002)	-0.001*** (0.0002)	-0.001*** (0.0002)	-0.001*** (0.0003)
<i>Dependent Mean</i>	0.7	11.2	11.7	6.1	8.1	17.2
<i>Observations</i>	499,440	499,440	499,440	499,440	499,440	499,440

Notes: Data from the LEHD Origin-Destination Employment Statistics (LODES), LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2002 to 2012. LODES data is only available starting in 2002. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For comparison, circulation of the *USA TODAY* was 4.5 percent at the 95th percentile. Earnings and Industry classifications are pre-defined in the LODES data. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A5

Sensitivity to Various Fixed Effects

	Inverse Hyperbolic Sine of the Number of Migrants $_{oSt}$					Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs $_{oSt}$				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>National Newspaper Exposure</i> $_{oSt}$	0.024*** (0.003)	0.025*** (0.004)	0.026*** (0.003)	0.026*** (0.004)	0.027*** (0.004)	0.112*** (0.007)	0.068*** (0.009)	0.060*** (0.006)	0.040*** (0.007)	0.024*** (0.006)
<i>National Newspaper Exposure</i> $^2_{oSt}$	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.004*** (0.0004)	-0.002*** (0.0004)	-0.002*** (0.0003)	-0.001*** (0.0003)	-0.001*** (0.0002)
<i>Origin by Destination F.E.</i>	X	X	X	X	X	X	X	X	X	X
<i>Destination by Year F.E.</i>		X		X	X		X		X	X
<i>Origin by Year F.E.</i>			X	X	X			X	X	X
<i>Origin State by Destination by Year</i>					X					X
<i>Dependent Mean (in Levels)</i>	7.6	7.6	7.6	7.6	7.6	31.4	31.4	31.4	31.4	31.4
<i>Observations</i>	590,224	590,224	590,224	590,224	590,224	499,440	499,440	499,440	499,440	499,440

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis, and the Alliance for Audited Media. The level of observation is the origin county by destination state by year. *National Newspaper Exposure* $_{oSt}$ is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. The inverse hyperbolic sine approximates a natural log transformation, but is defined for values of zero. In all specifications origin/destination pair fixed effects are included to control for time-invariant differences across pairs. Columns (5) and (10) include origin state by destination by year fixed effects, thus exploiting variation across origin counties in the same state to the same destination state. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A6

Sensitivity to Functional Form

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}				Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}			
	Linear (1)	Quadratic (2)	Cubic (3)	IHS (4)	Linear (5)	Quadratic (6)	Cubic (7)	IHS (8)
<i>National Newspaper Exposure_{oSt}</i>	0.014*** (0.003)	0.025*** (0.004)	0.022*** (0.005)		0.037*** (0.007)	0.068*** (0.009)	0.099*** (0.011)	
<i>National Newspaper Exposure_{oSt}²</i>		-0.001*** (0.0001)	-0.0002 (0.001)			-0.002*** (0.0004)	-0.006*** (0.001)	
<i>National Newspaper Exposure_{oSt}³</i>			0.00001 (0.00001)				0.0001*** (0.00002)	
<i>Inverse Hyperbolic Sine of National Newspaper Exposure_{oSt}</i>				0.046*** (0.007)				0.155*** (0.019)
<i>Dependent Mean (in Levels)</i>	7.6	7.6	7.6	7.6	31.4	31.4	31.4	31.4
<i>Observations</i>	590,224	590,224	590,224	590,224	499,440	499,440	499,440	499,440

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis, and the Alliance for Audited Media. The level of observation is the origin county by destination state by year. *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. The inverse hyperbolic sine approximates a natural log transformation, but is defined for values of zero. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A7

Sensitivity to Sample

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}					Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}				
	Include NYC DMA (1)	Include Fracking Origins (2)	Exclude Non-Fracking Origins in Fracking States (3)	Exclude Zero Circulation (4)	Exclude Top One Percent of Exposure (5)	Include NYC DMA (6)	Include Fracking Origins (7)	Exclude Non-Fracking Origins Fracking States (8)	Exclude Zero Circulation (9)	Exclude Top One Percent of Exposure (10)
<i>National Newspaper Exposure_{oSt}</i>	0.013*** (0.004)	0.025*** (0.003)	0.026*** (0.004)	0.024*** (0.004)	0.31*** (0.006)	0.045*** (0.006)	0.067*** (0.009)	0.066*** (0.009)	0.049*** (0.011)	0.123*** (0.013)
<i>National Newspaper Exposure_{oSt}²</i>	-0.0001** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0002)	-0.001*** (0.0001)	-0.0003 (0.001)	-0.0004*** (0.0001)	-0.002*** (0.0003)	-0.002*** (0.0004)	-0.001** (0.0004)	-0.007*** (0.001)
<i>Dependent Mean (in Levels)</i>	7.7	12.2	2.1	11.1	6.1	31.2	62.2	1.3	43.5	26.9
<i>Observations</i>	596,256	639,840	398,944	392,832	505,664	504,544	541,504	337,568	332,416	427,840

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis, and the Alliance for Audited Media. The level of observation is the origin county by destination state by year. *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A8

Impact of Destination State Specific National Newspaper Exposure on Migration, Restrict Sample to Treatment Years

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}				Number of Migrants _{oSt}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sample: 2009-2012								
<i>National Newspaper Exposure</i> _{oSt}	0.018*** (0.003)	0.018*** (0.003)	0.015*** (0.003)	0.015*** (0.003)	1.416*** (0.372)	1.441*** (0.379)	1.314*** (0.427)	1.315*** (0.422)
<i>National Newspaper Exposure</i> _{oSt} ²	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.034** (0.020)	-0.035** (0.020)	-0.027* (0.022)	-0.031** (0.021)
<i>Mean Number of Migrants</i>	8.2	8.2	8.2	8.2	8.2	8.2	8.2	8.2
<i>Observations</i>	181,696	181,696	181,696	181,696	181,696	181,696	181,696	181,696
Sample: 2011-2012								
<i>National Newspaper Exposure</i> _{oSt}	0.014*** (0.004)	0.014*** (0.004)	0.013*** (0.005)	0.013*** (0.005)	0.358*** (0.128)	0.365*** (0.129)	0.219* (0.126)	0.216* (0.129)
<i>National Newspaper Exposure</i> _{oSt} ²	-0.001*** (0.0003)	-0.001*** (0.0003)	-0.001*** (0.0003)	-0.001*** (0.0003)	-0.016*** (0.006)	-0.017*** (0.006)	-0.013** (0.006)	-0.015** (0.006)
<i>Mean Number of Migrants</i>	9.1	9.1	9.1	9.1	9.1	9.1	9.1	9.1
<i>Observations</i>	90,848	90,848	90,848	90,848	90,848	90,848	90,848	90,848

Notes: Data from the IRS Statistics of Income, LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable *National Newspaper Exposure*_{oSt} is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For origin/destination pairs with any news exposure, mean national newspaper exposure is 0.99. Each column replicates the corresponding estimate from Table 2, but restricts the years in the sample. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A9

Accounting for Censoring: Impact of Destination State Specific Newspaper Exposure on Migration

	Inverse Hyperbolic Sine of the Number of Migrating Tax Units _{oSt} Lower Bound: As Reported Replace 0 with 9		Number of Migrating Tax Units _{oSt} Lower Bound: As Reported Replace 0 with 9		Over 10 Migrating Tax Units _{oSt}	Inverse Hyperbolic Sine of the Number of Migrants _{oSt} Positive Flows in All Years
	(1)	(2)	(3)	(4)	(5)	(6)
<i>National Newspaper Exposure_{oSt}</i>	0.022*** (0.003)	0.001*** (0.0003)	0.838*** (0.222)	0.732*** (0.209)	0.005*** (0.001)	0.018** (0.007)
<i>National Newspaper Exposure_{oSt}²</i>	-0.001*** (0.0001)	-0.00002** (0.00001)	-0.019** (0.010)	-0.016* (0.009)	-0.0002*** (0.00003)	-0.0004** (0.0002)
<i>Dependent Mean (in Levels)</i>	4	145	4	145	0.03	348.7
<i>Observations</i>	590,224	590,224	590,224	590,224	590,224	12,092

Notes: Data obtained from the IRS Statistics of Income, LexisNexis Newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012. The variable *National Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For comparison, circulation of the *USA TODAY* was 4.5 percent at the 95th percentile. In Columns (1) and (2) the outcome is the inverse hyperbolic sine of migrating tax units (rather than migrants). Censored values are assigned a value of 0 in Column (1), and assigned a value of 9 in Column (2), to provide a lower bound. In Columns (3) and (4) the outcome is the number of migrating tax units in levels, to account for the fact that percentages are not comparable when censored values are reassigned a value of 9. The outcome in Column (5) is an indicator that equals one if there were over 10 migrating tax units. During the sample period, flows with less than 10 returns were censored, and this outcome captures transitions across the censoring threshold. The outcome in Column (6) is the inverse hyperbolic sine of migrants for a subsample of origin/destination pairs that reported positive flows in all years. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. Estimates are similar if I instead impute missing migration flows as the average or maximum non-censored origin/destination flows rather than zero. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A10

Impact of Destination State Specific National Newspaper Exposure on Migration Rates to Fracking Counties in State

	Inverse Hyperbolic Sine of the Migration Rate per 100,000 _{oSt}				Migration Rate per 100,000 _{oSt}			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>National Newspaper Exposure</i> _{oSt}	0.020*** (0.003)	0.020*** (0.003)	0.023*** (0.003)	0.022*** (0.003)	0.400* (0.222)	0.438* (0.235)	0.663** (0.260)	0.622** (0.290)
<i>National Newspaper Exposure</i> _{oSt} ²	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.018 (0.011)	-0.019* (0.011)	-0.024** (0.012)	-0.025** (0.011)
<i>Local Newspaper Exposure</i> _{oSt}				0.013** (0.005)				-0.629 (2.617)
<i>Local Newspaper Exposure</i> _{oSt} ²				-0.0001** (0.00005)				0.063 (0.056)
<i>Origin Labor Market Controls</i>		X				X		
<i>Origin by Year Effects</i>			X	X			X	X
<i>Origin/Destination Local News</i>				X				X
<i>Mean Migrants per 100,000</i>	6.0	6.0	6.0	6.0	6.0	6.0	6.0	6.0
<i>Observations</i>	590,224	590,224	590,224	590,224	590,224	590,224	590,224	590,224

Notes: Data from the IRS Statistics of Income, LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year from 2000 to 2012. The outcome is the number of migrants divided by the origin county population in 2000, in hundreds of thousands of people. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable *National Newspaper Exposure*_{oSt} is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For origin/destination pairs with any news exposure, mean national newspaper exposure is 0.99. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Origin controls include the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$). Origin/destination specific local news is all destination state specific fracking news content listed in LexisNexis from non-national domestic newspapers. The variable *Local Newspaper Exposure*_{oSt} is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 40 percent circulation rate, approximately the 95th percentile of pre-fracking circulation among non-national newspapers with articles about fracking. The sample correlation between national news exposure and local news exposure is approximately 0.12. Origin county by year fixed effects control for time-varying characteristics of the origin county and account for potential changes in preferences toward fracking that might be correlated with newspaper readership and affect migration to fracking areas. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A11

Demographic Characteristics of *USA TODAY* Readership in 2007

	Readership Characteristics (1)	Fraction Adult Population (2)
<i>Total Readers</i>	3,882,000	0.017
<i>Male</i>	0.68	0.024
<i>Female</i>	0.32	0.011
<i>18-49</i>	0.63	0.018
<i>25-54</i>	0.66	0.20
<i>35-54</i>	0.50	0.022
<i>40-59</i>	0.48	0.022
<i>No College</i>	0.29	0.009
<i>Some College (<4-year)</i>	0.28	0.022
<i>4-year Degree or More</i>	0.43	0.029
<i>Professional/Managerial</i>	0.34	0.084
<i>Top/Middle Manager</i>	0.25	—
<i>Employed</i>	0.80	0.022
<i>HH Income <50K</i>	0.31	0.013
<i>HH Income 50-75K</i>	0.17	0.005
<i>HH Income 75-100K</i>	0.14	0.018
<i>HH Income ≥100K</i>	0.38	0.007
<i>No Children</i>	0.55	0.015
<i>Any Children</i>	0.45	0.020
<i>Homeowners</i>	0.74	0.018
<i>Renters</i>	0.26	0.016

Notes: Readership characteristics based on the MRI Fall 2007 Report (http://usatoday30.usatoday.com/media_kit/usatoday/au_general_demographics.htm). Age groups are as defined in the 2007 Report and are not mutually exclusive groups. Readers' education, income, presence of children, and homeownership status are grouped into mutually exclusive groups based on the reported groups from the 2007 report. The population measures in Column (2) are constructed using the ACS 2007 microdata for the full US population 18 and older. Top and middle managers are not identified in the ACS.

Table A12

County Characteristics by Local Newspaper Circulation Among Counties at the Market Border

	Characteristic in 2000		Change from 2000 to 2010		Within Border Change from 2000 to 2010
	In-Market (1)	Out of Market (2)	In-Market (3)	Out of Market (4)	
<i>Migrants to fracking areas (Percent of Population)</i>	0.22	0.24	-0.018	-0.021	0.00
<i>Employment to Population (16+)</i>	57.92	57.77	-0.80	-0.96	0.161*
<i>Unemployment Rate</i>	3.44	3.40	1.07	1.13	-0.06
<i>Median Household Income</i>	37,695	37,599	9,358	9,562	-185
<i>Percent in Poverty</i>	12.87	12.86	1.46	1.31	0.15*
<i>Percent White</i>	84.07	85.17	-2.08	-1.77	-0.31**
<i>Percent Black</i>	8.42	7.62	0.27	0.23	0.06
<i>Percent Hispanic</i>	6.75	6.90	2.44	2.28	0.20***
<i>Percent Other Race</i>	7.51	7.21	1.81	1.54	0.26***
<i>Percent Population 20-34</i>	18.78	18.08	-0.48	-0.70	0.20***
<i>Percent Population 35-64</i>	38.74	39.40	1.66	1.69	-0.04
<i>Percent Population Over 64</i>	14.54	14.44	0.95	1.30	-0.33***
<i>Percent Households Renting</i>	27.66	25.85	1.63	1.59	0.05

Notes: Migration data from the IRS Statistics of Income. Other county characteristics obtained through American FactFinder from the 2000 Census and 2010 Census and 5-Year American Community Survey. Local newspaper circulation data from the Alliance for Audited Media. Columns (1) and (2) report the means for counties inside versus outside the local newspapers' distribution markets. Columns (3) and (4) reports the average change from 2000 to 2010 for these groups. Column (5) reports the differences between Columns (3) and (4) when controlling for newspaper border, to compare counties in the same newspaper border. Standard errors corrected for clustering at the origin DMA level. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *.

Table A13

Advertising Effects of News Exposure: Market Expanding or Share Stealing

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}		Inverse Hyperbolic Sine of the Number of Cross-County Commute Jobs _{oSt}	
	(1)	(2)	(3)	(4)
<i>National Newspaper Exposure_{oSt}</i>	0.027*** (0.004)	0.027*** (0.003)	0.041*** (0.007)	0.049*** (0.007)
<i>National Newspaper Exposure_{oSt}²</i>	-0.001*** (0.0001)	-0.001*** (0.0001)	-0.001*** (0.0003)	-0.001*** (0.0003)
<i>All States Newspaper Exposure_{ot}</i>	0.0003 (0.0003)		0.005*** (0.001)	
<i>All States Newspaper Exposure_{ot}²</i>	-0.00001** (0.000003)		-0.00003*** (0.00001)	
<i>Max. State Newspaper Exposure_{ot}</i>		0.001 (0.001)		0.015*** (0.004)
<i>Max. State Newspaper Exposure_{ot}²</i>		-0.0001** (0.0001)		-0.0004*** (0.0002)
<i>Dependent Mean (in Levels)</i>	7.6	7.6	31.4	31.4
<i>Observations</i>	590,224	590,224	499,440	499,440

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis, and the Alliance for Audited Media. The level of observation is the origin county by destination state by year. *Newspaper Exposure_{oSt}* is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. *All States Newspaper Exposure_{ot}* is the total news exposure for all 16 destination states within an origin year, to determine if news about fracking in general affects migration. *Max. States Newspaper Exposure_{ot}* is the highest level of news exposure across all 16 destination state within an origin year, to determine if higher news exposure leads to shifting away from other fracking destinations. Controls for the origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A14

Chain Migration: Impact of Lagged Newspaper Exposure on Migration

	Inverse Hyperbolic Sine of the Number of Migrants _{oSt}					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>National Newspaper Exposure</i> _{oSt}	0.025*** (0.004)	0.022*** (0.003)	0.022*** (0.004)	0.024*** (0.004)	0.022*** (0.003)	0.022*** (0.004)
<i>National Newspaper Exposure</i> _{oSt} ²	-0.001*** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)	-0.001** (0.000)	-0.001*** (0.000)	-0.001*** (0.000)
<i>National Newspaper Exposure</i> _{oSt-1}		0.007* (0.004)	0.008** (0.004)		0.007* (0.004)	0.008** (0.004)
<i>National Newspaper Exposure</i> _{oSt-1} ²		-0.000 (0.000)	-0.000 (0.000)		-0.000 (0.000)	-0.000 (0.000)
<i>National Newspaper Exposure</i> _{oSt-2}			-0.003 (0.006)			-0.004 (0.006)
<i>National Newspaper Exposure</i> _{oSt-2} ²			0.001 (0.001)			0.001 (0.001)
<i>IHS of Number of Migrants</i> _{oSt-1}				0.042*** (0.012)	0.042*** (0.012)	0.032** (0.012)
<i>Dependent Mean</i>	7.6	7.7	7.7	7.7	7.7	7.7
<i>Observations</i>	590,224	544,832	499,440	544,832	544,832	499,440

Notes: Data from the IRS Statistics of Income, LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year for the years specified. Origin counties with any fracking production or in the New York City designated market area are excluded. The variable National Newspaper Exposure is scaled such that a one unit increase represents the impact of one additional news story in a newspaper with a 5 percent circulation rate. For origin/destination pairs with any news exposure, mean national newspaper exposure is 0.99. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Origin county by year fixed effects control for time-varying characteristics of the origin county and account for potential changes in preferences toward fracking that might be correlated with newspaper readership and affect migration to fracking areas. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A15

Reverse Migration: Impact of National Newspaper Exposure at Origin on Future Migration from Destination

	Inverse Hyperbolic Sine of the Number of Migrants from S to o			
	(1)	(2)	(3)	(4)
<i>National Newspaper Exposure</i> _{oSt}	0.005 (0.003)	0.004 (0.003)	0.003 (0.003)	0.004 (0.003)
<i>National Newspaper Exposure</i> _{oSt} ²	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
<i>National Newspaper Exposure</i> _{$oSt-1$}		0.008* (0.005)	0.006 (0.005)	0.007 (0.005)
<i>National Newspaper Exposure</i> _{$oSt-1$} ²		-0.000 (0.000)	-0.000 (0.000)	-0.000 (0.000)
<i>National Newspaper Exposure</i> _{$oSt-2$}			0.016** (0.007)	0.017** (0.007)
<i>National Newspaper Exposure</i> _{$oSt-2$} ²			-0.003** (0.001)	-0.003** (0.001)
<i>National Newspaper Exposure</i> _{$oSt-3$}				-0.012 (0.026)
<i>National Newspaper Exposure</i> _{$oSt-3$} ²				-0.012 (0.017)
<i>Dependent Mean</i>	7.4	7.4	7.4	7.5
<i>Observations</i>	590,224	544,832	499,440	454,016

Notes: Data from the IRS Statistics of Income, LexisNexis newspaper transcripts, and newspaper circulation from the Alliance for Audited Media. The level of observation is the origin county by destination state by year for the years specified. Origin counties with any fracking production or in the New York City designated market area are excluded. The outcome is the number of migrants from fracking counties in state S to origin county o . National Newspaper Exposure is exposure at the origin about the destination. This specification will capture how news exposure in previous years affects reverse migration flows. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *

Table A16

Heterogeneous Impacts by Origin Employment to Population Ratio

	Inverse Hyperbolic Sine		Levels	
	Number of Migrants (1)	Number of Commuters (2)	Number of Migrants (3)	Number of Commuters (4)
<i>Newspaper Exposure</i> _{oSt}	0.0270*** (0.0037)	0.0731*** (0.0088)	1.7029*** (0.5103)	5.3651** (2.2426)
<i>Newspaper Exposure</i> _{oSt} ²	-0.0007*** (0.0001)	-0.0020*** (0.0004)	-0.0392* (0.0234)	-0.2425** (0.1048)
<i>Newspaper Exposure</i> _{oSt} *	-0.0006*** (0.0002)	-0.0017** (0.0006)	-0.0100 (0.0221)	-0.1350 (0.0958)
<i>Emp/Pop</i> _{ot-1}				
<i>Newspaper Exposure</i> _{oSt} *	-0.0001*** (0.000)	-0.0002*** (0.000)	-0.0040*** (0.0013)	-0.0189*** (0.0066)
<i>Emp/Pop</i> _{ot-1} ²				
<i>Newspaper Exposure</i> _{oSt} ² *	0.000 (0.000)	0.000 (0.0001)	-0.0009 (0.0011)	0.0058 (0.0043)
<i>Emp/Pop</i> _{ot-1}				
<i>Newspaper Exposure</i> _{oSt} ² *	-0.000 (0.000)	0.000 (0.000)	-0.000 (0.0001)	0.0008** (0.0004)
<i>Emp/Pop</i> _{ot-1} ²				
<i>Emp/Pop</i> _{ot-1}	0.0004** (0.0002)	0.0002 (0.0006)	0.0052 (0.0192)	-0.1376 (0.1236)
<i>Emp/Pop</i> _{ot-1} ²				
<i>Emp/Pop</i> _{ot-1}	0.000 (0.000)	-0.000 (0.000)	0.0006* (0.0004)	0.0032 (0.0021)
<i>Dependent Mean (in levels)</i>	7.635	31.23	7.635	31.23
<i>Observations</i>	544,688	499,296	544,688	499,296

Notes: Data obtained from the IRS Statistics of Income, LEHD Origin-Destination Employment Statistics (LODES), LexisNexis, and the Alliance for Audited Media. The level of observation is the origin county by destination state by year. The origin county employment to population ratio (Emp/Pop) is obtained from the BLS, and lagged by one year. Emp/Pop is demeaned, such that the direct effect of newspaper exposure is the effect for a county at the mean employment to population ratio (70.9 percent). Controls for the current origin county unemployment rate, employment to population ratio, and average annual earnings (2010\$) are also included. In all specifications origin/destination pair fixed effects, and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the designated market area to account for correlation across geography and time. p<0.01 ***, p<0.05 **, p<0.1 *.

Table A17

Simulated Impacts of News Exposure Migration Flows by Origin Employment to Population Ratio

	Employment to Population Ratio in $t - 1$ ($\mu = 70.9$)				
	60	65	70	75	80
<i>Exposure Level</i>					
<i>0</i>	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
<i>1</i>	0.027*** (0.003)	0.028*** (0.004)	0.027*** (0.004)	0.023*** (0.003)	0.017*** (0.003)
<i>5</i>	0.116*** (0.015)	0.125*** (0.016)	0.120*** (0.016)	0.103*** (0.015)	0.072*** (0.014)
<i>10</i>	0.189*** (0.025)	0.210*** (0.027)	0.205*** (0.027)	0.174*** (0.026)	0.116*** (0.026)
<i>15</i>	0.218*** (0.031)	0.256*** (0.034)	0.255*** (0.035)	0.214*** (0.035)	0.133*** (0.040)
<i>Mean Exposure in 2012</i>	1.57	2.08	2.48	2.44	1.89
<i>Implied Impact</i>	0.041*** (0.005)	0.057*** (0.007)	0.064*** (0.008)	0.054*** (0.008)	0.030*** (0.006)
<i>Maximizing Exposure</i>	15.9	18.4	19.6	18.9	15.6
<i>Implied Impact</i>	0.219*** (0.032)	0.265*** (0.038)	0.270*** (0.041)	0.223*** (0.043)	0.133*** (0.042)

Notes: Simulated impacts are obtained for each combination of origin employment to population ratio and exposure level from equation (7), where the outcome is the inverse hyperbolic sine of the number of migrants. The corresponding coefficients are reported in Appendix Table A11. The maximizing exposure is obtained by setting the first derivative of equation (7) with respect to newspaper exposure equal to zero and solving for the maximizing exposure for the specified employment to population ratio. This value is rounded down to the nearest whole number. The implied impact is the corresponding effect of the maximizing exposure level. For reference, the mean exposure level in 2012 for origin counties with employment to population ratios within 2.5 percent of the specified threshold. An employment to population ratio of 60 roughly corresponds to the 15th percentile while a ratio of 80 corresponds to roughly the 85th percentile. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *.

Table A18

Simulated Impacts of News Exposure Cross-County Commute Flows by Origin Employment to Population Ratio

	Employment to Population Ratio in $t - 1$ ($\mu = 70.9$)				
	60	65	70	75	80
<i>Exposure Level</i>					
0	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
1	0.071*** (0.010)	0.076*** (0.009)	0.073*** (0.009)	0.062*** (0.009)	0.043*** (0.011)
5	0.315*** (0.041)	0.336*** (0.040)	0.322*** (0.039)	0.274*** (0.040)	0.191*** (0.046)
10	0.527*** (0.70)	0.566*** (0.071)	0.544*** (0.071)	0.461*** (0.070)	0.319*** (0.074)
15	0.637*** (0.089)	0.689*** (0.099)	0.664*** (0.102)	0.562*** (0.098)	0.384*** (0.093)
<i>Mean Exposure in 2012</i>	1.57	2.08	2.48	2.44	1.89
<i>Implied Impact</i>	0.110*** (0.015)	0.153*** (0.018)	0.172*** (0.020)	0.144*** (0.021)	0.079*** (0.020)
<i>Maximizing Exposure</i>	17.9	18.3	18.5	18.4	17.7
<i>Implied Impact</i>	0.654*** (0.099)	0.713*** (0.118)	0.689*** (0.128)	0.582*** (0.120)	0.393*** (0.105)

Notes: Simulated impacts are obtained for each combination of origin employment to population ratio and exposure level from equation (7), where the outcome is the inverse hyperbolic sine of the number of cross-county commuters. The corresponding coefficients are reported in Appendix Table A11. The maximizing exposure is obtained by setting the first derivative of equation (7) with respect to newspaper exposure equal to zero and solving for the maximizing exposure for the specified employment to population ratio. This value is rounded down to the nearest whole number. The implied impact is the corresponding effect of the maximizing exposure level. For reference, the mean exposure level in 2012 for origin counties with employment to population ratios within 2.5 percent of the specified threshold. An employment to population ratio of 60 roughly corresponds to the 15th percentile while a ratio of 80 corresponds to roughly the 85th percentile. $p < 0.01$ ***, $p < 0.05$ **, $p < 0.1$ *.

Table A19

State-level Events Controlled for in Google Trends State Specifications

Date (1)	Event (2)	States (3)
November 7-9, 2006	Four-way Texas Gubernatorial Election	Texas
September 9-10, 2010	San Bruno Pipeline Explosion	California
September 6-11, 2011	2011 Texas Wildfires	Texas
Dec. 31, 2011-Jan. 2, 2012	4.0 Earthquake in Eastern Ohio	Ohio
February 19-20, 2012	Texas A&M v. Oklahoma State Basketball Game	Texas, Oklahoma
February 27-28, 2012	Chardon High School shooting	Ohio
March 7, 2012	Ohio Primary Elections	Ohio
May 21-22, 2012	Arkansas Primary Elections	Arkansas

Notes: High-level interest events that are closely tied to a specific state during the Google Trend search windows are controlled for to increase precision. Indicators that equal one for each of the listed dates for the destination state listed are included in the state name Google Trend analysis.

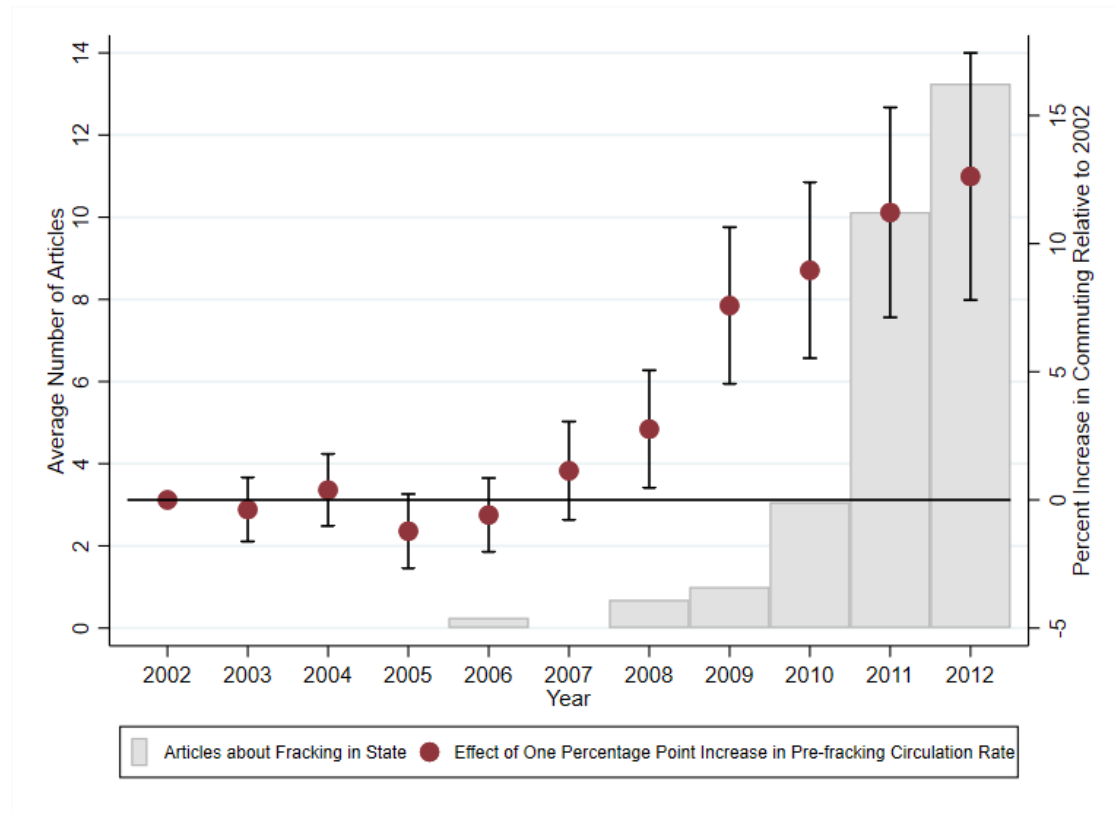


Figure A1

Trends in Commuting by Pre-fracking Circulation

Notes: For each origin, the pre-fracking circulation rate is the weighted average of the pre-fracking circulation of the *USA TODAY*, *New York Times*, and *Wall Street Journal*, where weights are the share of the total articles about fracking in each newspaper. This measure is then interacted with year indicators. Commuting data is only available starting in 2002, so 2002 is treated as the base year. The inverse hyperbolic sine of the number of cross-county commuting jobs is then regressed on this set of interactions along with origin-destination pair effects and destination-by-year fixed effects, as in the main specification, to trace out the effect of a one percentage point increase in the pre-fracking circulation rate on migration, as a percent. The marginal effect of one unit of a one percentage point increase is converted to percentage points and plotted for each year on the right axis, to look at trends by pre-fracking circulation. Standard errors from the regressions are corrected for clustering at the origin DMA level. For reference, the average number of articles about fracking in each state is also plotted for each year in bars on the left axis.

Source: Author's calculations using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and commuting flows from the LODES.




Figure A2

Marginal Impact of Newspaper Exposure on Migration by Origin to Destination Distance

Notes: Coefficients and confidence intervals plotted for the marginal effect of newspaper exposure on migration flows from equation (4), estimated over one hundred mile bins. Standard errors calculated using the Delta Method.

Source: Author's calculations using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and migration flows from the IRS SOI.

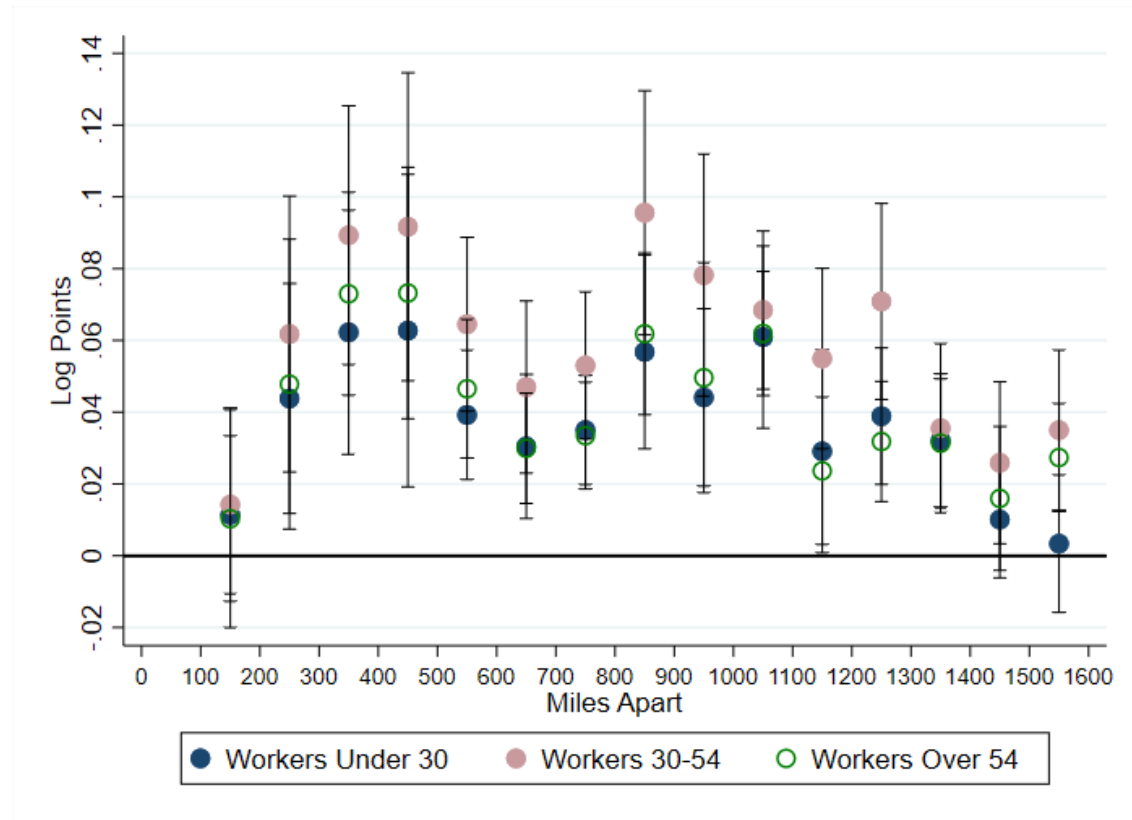


Figure A3

Marginal Impact of Newspaper Exposure on Commuting by Age and Origin to Destination Distance

Notes: Coefficients and confidence intervals plotted for the marginal effect of newspaper exposure on cross-county commute flows from equation (13), estimated over one hundred mile bins, separately by worker age. Standard errors calculated using the Delta Method. The estimated marginal impact for pairs less than one hundred miles apart for each age group are highly negative, at -0.06 (0.04), -0.10 (0.06), and -0.16 (0.06), respectively.

Source: Author's calculations using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and commute flows from the LEHD Origin-Destination Employment Statistics (LODES).

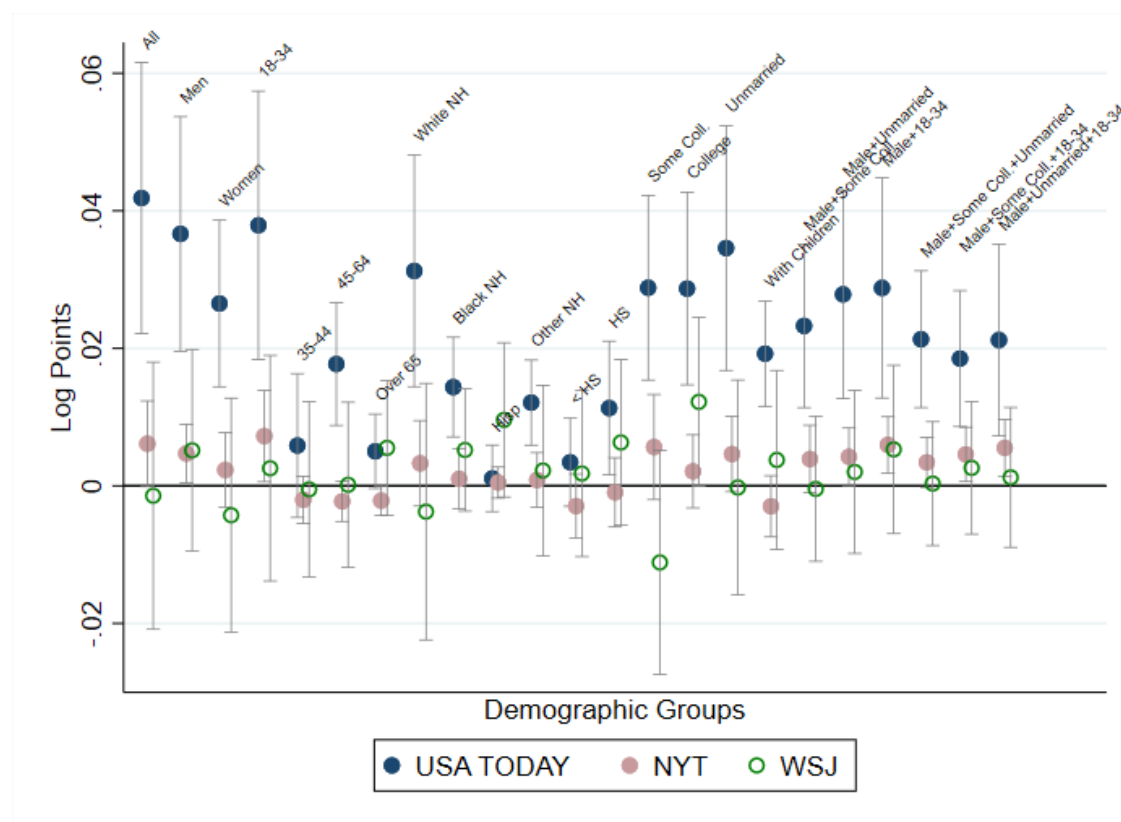


Figure A4

Heterogeneity by Newspaper: Migration Impact of Newspaper Exposure by Demographic Group

Notes: The total impact of a one unit increase in exposure for each of the three national newspapers is plotted with 95 percent confidence intervals. The level of observation is the origin MIGPUMA by destination state by year from 2005 to 2012 for the given group. Origin MIGPUMAs with any fracking production are excluded. Each newspaper's exposure level is scaled to represent the impact of one additional news story in a county with circulation at the 95th percentile. Origin/destination pair fixed effects and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the origin state.

Source: Author's calculation using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and microdata from the 2005-2012 American Community Survey.




Figure A5

Household Chain Migration in ACS: Migration Impact of Newspaper Exposure on Migration into Pre-existing Households

Notes: The total impact of a one unit increase in exposure for each of the three national newspapers is plotted with 95 percent confidence intervals. The level of observation is the origin MIGPUMA by destination state by year from 2005 to 2012 for the given group. Origin MIGPUMAs with any fracking production are excluded. Each newspaper's exposure level is scaled to represent the impact of one additional news story in a county with circulation at the 95th percentile. Origin/destination pair fixed effects and destination by year fixed effects are included to control for time-invariant differences across pairs and characteristics of the destination that vary over time. Standard errors are adjusted for clustering at the origin state.

Source: Author's calculation using circulation rates from the Alliance of Audited Media, newspaper content from LexisNexis, and microdata from the 2005-2012 American Community Survey.

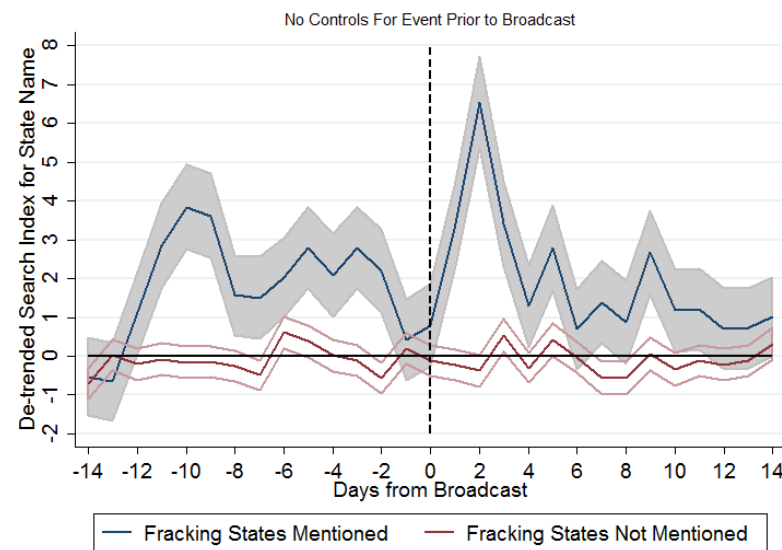
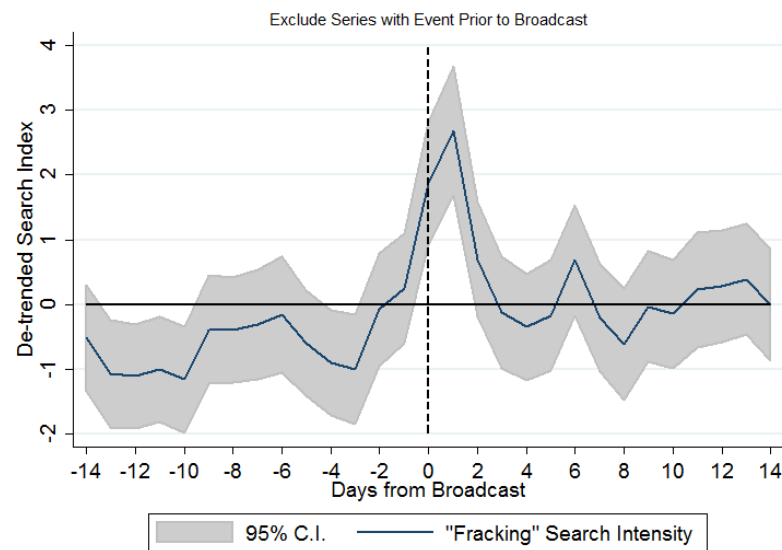


Figure A6

Google Search Interest: Exclude Event Controls

Notes: Plot depicts the same average daily search index for “fracking” and specific states mentioned as in Figures VI and VII, but does not include controls for high publicity events that occurred during the search window and were either related to fracking, or a specific state. Standard errors are clustered at the search level.

Source: Author’s calculations using daily search indices from Google Trends.

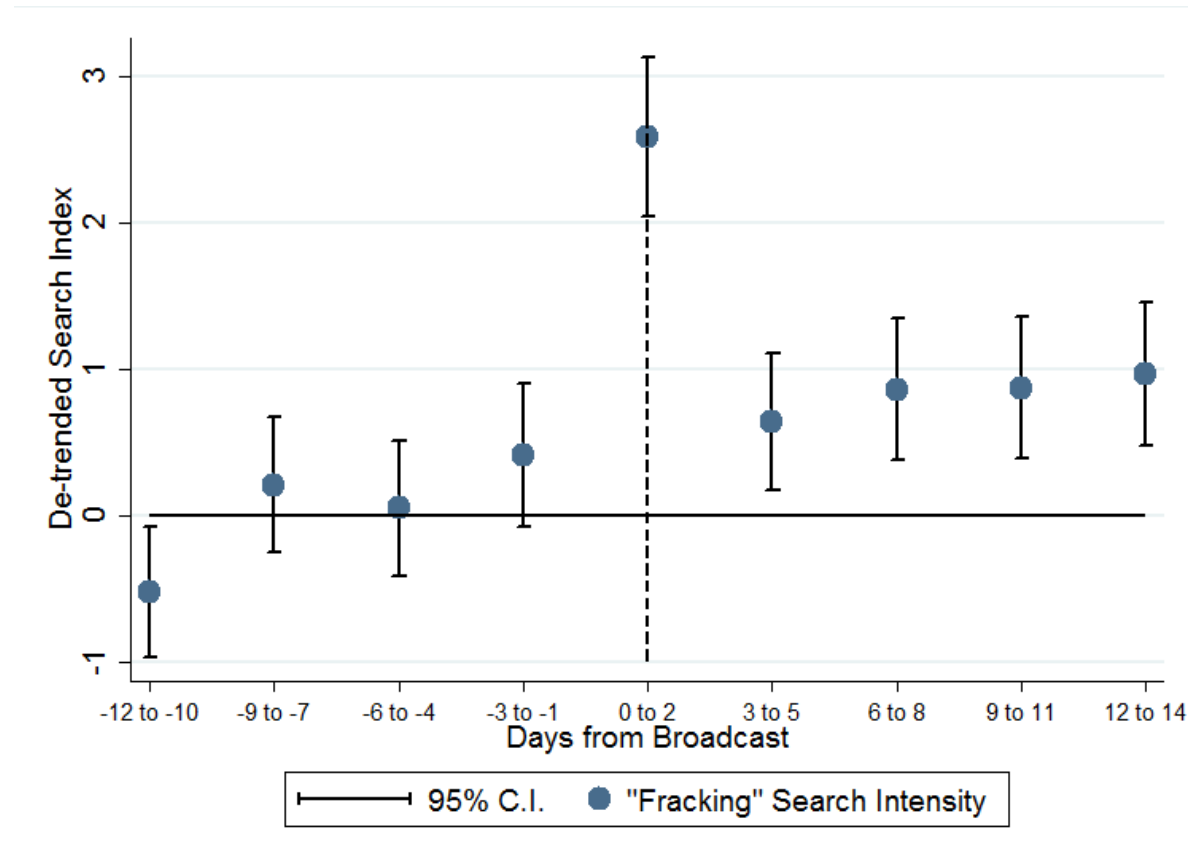


Figure A7

Google Search Interest in “Fracking” After TV News Broadcasts, 3 Day Bins

Notes: Plot depicts the average daily search index for the term “fracking” by DMA before and after 17 TV broadcast mentioning fracking or shale gas between 2006 and 2012 as recorded by the Vanderbilt Television News Archive, as in Figure VIII, but groups days into 3-day bins. Search intensity is de-trended by removing day of week and search (DMA by four week publication window) specific effects. To be consistent with other analysis in the paper, one broadcast from CNN and one broadcast from Fox News are excluded. Four days prior to a news broadcast on January 28, 2012, President Barack Obama mentioned shale gas exploration due to fracking in the State of the Union Address. Four days prior to a news broadcast on January 4, 2012, there was an earthquake in Ohio that reporters linked to fracking. For both of these event I include indicator variables for the next four days. When these events are not controlled for, there is a marginally significant increase in search intensity in the days prior to the broadcast. Standard errors are clustered at the search level.

Source: Author’s calculations using daily search indices from Google Trends.

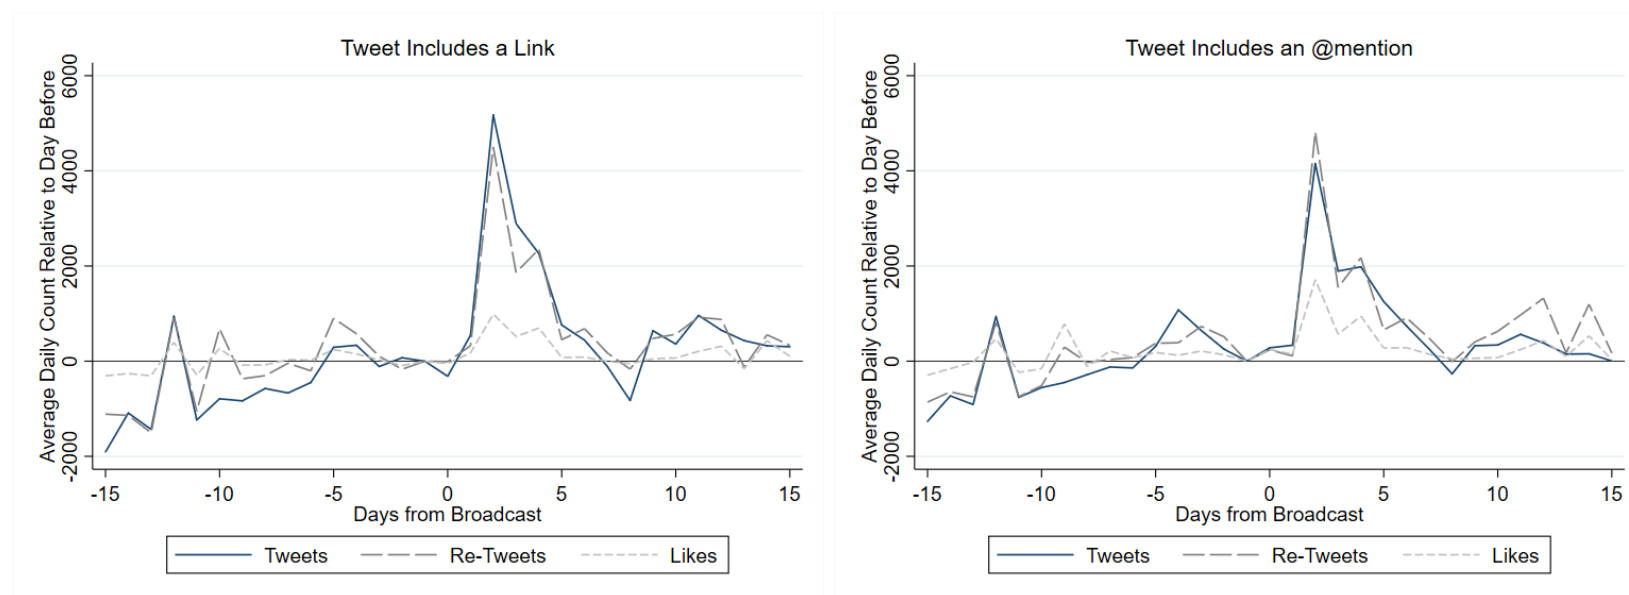


Figure A8

Average Tweets, Re-tweets, and Likes of Tweets about Fracking that Include Links or @ Mentions

Notes: All tweets with the word “fracking” and either a link or an @ mention within 15 days of the post-2009 TV news broadcast events are included. The daily number of total tweets, re-tweets, and likes are averaged over all of the broadcast events. The Twitter-user’s location is not attainable, and I am unable to exploit geographic variation in exposure to news about fracking. Links might direct people to other websites with more information.

Source: Author’s calculations using tweets collected from Twitter.

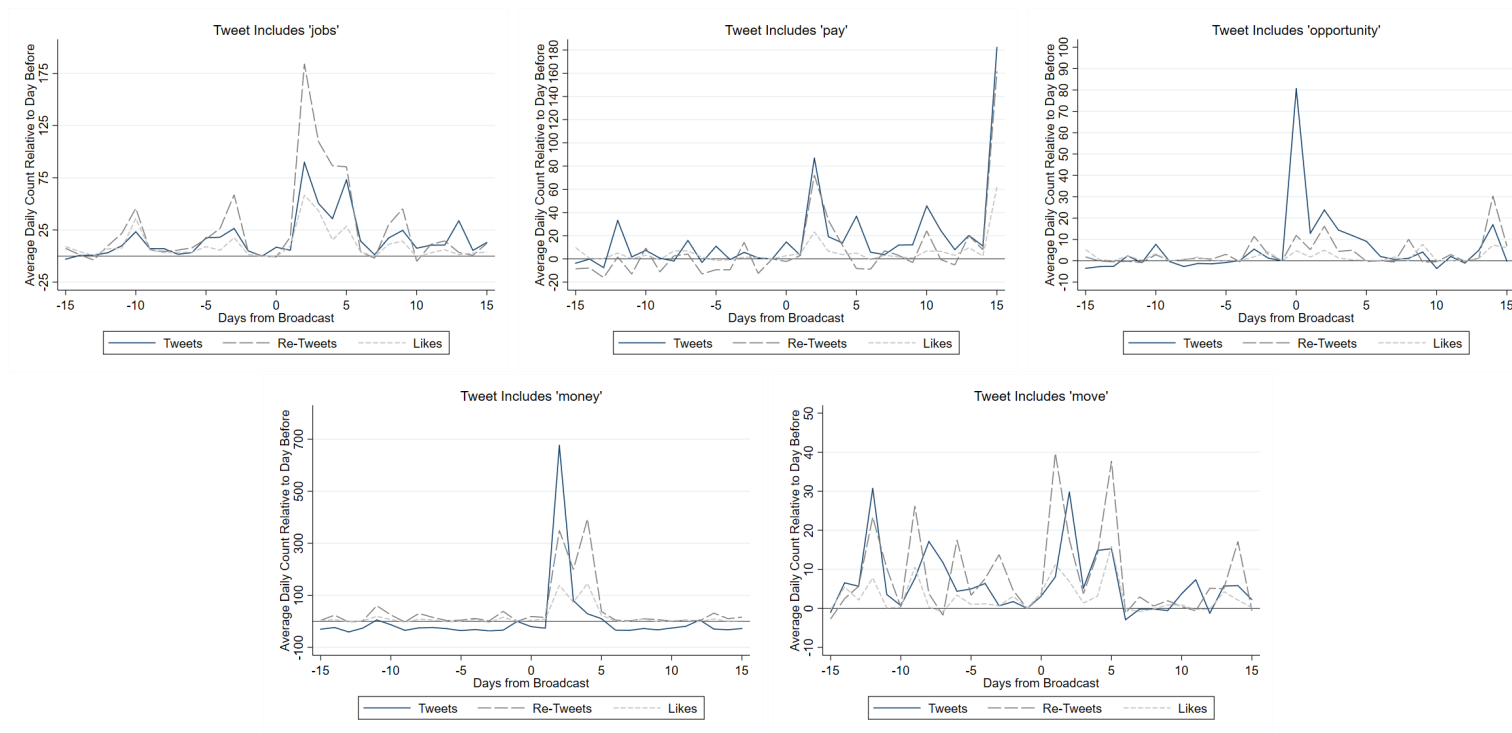


Figure A9

Average Tweets, Re-tweets, and Likes of Tweets about Fracking that Include Job Related Terms

Notes: All tweets with the word “fracking” and the listed word (i.e., “jobs”) within 15 days of the post-2009 TV news broadcast events are included. The daily number of total tweets, re-tweets, and likes are averaged over all of the broadcast events. The Twitter-user’s location is not attainable, and I am unable to exploit geographic variation in exposure to news about fracking.

Source: Author’s calculations using tweets collected from Twitter.

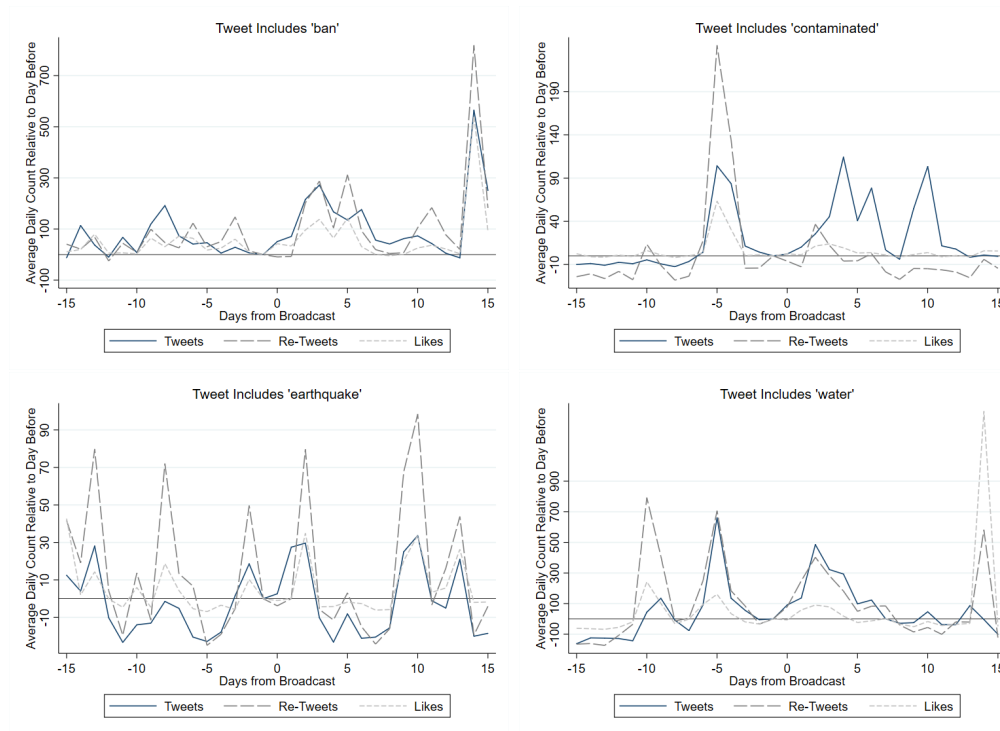


Figure A10

Average Tweets, Re-tweets, and Likes of Tweets about Fracking that Include Environmental Related Terms

Notes: All tweets with the word “fracking” and the listed word (i.e., “ban”) within 15 days of the post-2009 TV news broadcast events are included. The daily number of total tweets, re-tweets, and likes are averaged over all of the broadcast events. The Twitter-user’s location is not attainable, and I am unable to exploit geographic variation in exposure to news about fracking.

Source: Author’s calculations using tweets collected from Twitter.

For Online Publication: Appendix B. Supplemental Data Appendix

Below I describe each of the key datasets used in my analysis, as well as important characteristics of data construction.

DrillingInfo Oil and Gas Production Data

Well level information on drilling date, lease agreements, location, direction, and geological formation as well as other characteristics are provided through a restricted use data agreement from DrillingInfo. This data is proprietary, and obtained through an academic use agreement with DrillingInfo, available through their academic outreach initiative. DrillingInfo does not indicate if a well is a fracking well, as fracking is a means of stimulating production. To infer wells that are affected by the technological innovation associated with fracking, I use details on drilling direction and well location. Localized fracking booms occurred in part because of the combination of horizontal (directional) drilling and hydraulic fracturing. The DrillingInfo data reports whether a well is horizontally or vertically drilled. In addition, fracking was particularly impactful over shale plays, as these resources were not extractable previously. For this reason I assign non-vertical wells drilled in counties that intersect with shale plays as fracking wells. This production data is then combined with shale play boundary shapefiles provided by the Energy Information Administration to identify wells in shale plays and counties with fracking production.

Internal Revenue Service Statistics of Income County Flows

The Internal Revenue Service (IRS) Statistics of Income (SOI) division provides annual counts of county-to-county flows. This provides the raw number of tax returns and exemptions that were filed in one county in year $t - 1$ and in another county in year t . Each year, the IRS provides county-to-county flows of exemptions in a file with two years (e.g., 2002to2003). This represents exemptions that were in one county when filing in 2002 and in another county when filing in 2003. As most people file in the beginning of the year before April, I assign this flow to the year 2002. Most filing occurs between February and April, so annual migration flows capture moves from approximately March or April from one year to the next. For example, migrants who moved between March/April of 2011 and March/April 2012 will be assigned the year 2011. This introduces a slight lag relative to the measurement of news (from January to December).

Using exemptions to approximate people in a household, I can identify origin-destination county level flows. For privacy purposes, the IRS suppresses county pairs that have fewer than ten returns move in each year. As such, county pairs that have small, positive flows will be recorded as zero. This potentially introduces measurement error. For this reason, I also consider lower bound specifications where all county to county flows of zero are replaced with nine. This operates under the assumption that all flows had at least nine households move, which is likely an extreme over-estimate.

In 2011, the IRS changed the methodology for constructing the migration flows. Prior to 2011, only returns filed before late September were included in the calculations. From 2011 on, returns filed through the end of December were included. This led to greater representation of high income households (Pierce, 2015). They also expanded the way that matches were identified to consider all heads, spouses, and dependents. Using both the new method and old method, the IRS calculated state-level net migration rates to determine how much the series was affected. They find that 44 of the states (plus DC) differed by less than 5 percent and only Wyoming varied by more than 10 percent. Geographic differences in how the methodological change affects the calculated flow is potentially problematic, but addressed in the specification. If flows into certain

destinations were more affects, this will be picked up in the destination by year fixed effects. If flows out of certain origins were more affects, this will be picked up when I include the origin by year fixed effects. As seen above, this does not seem to affect the estimates. As a further check, I exclude each destination state one-by-one to see if the effect is only driven by one state. It is not. Also, the effects are essentially the same if I exclude fracking destinations where the methodological change had the biggest affect on flows: Wyoming, Louisiana, and Mississippi. In 2013 the suppression threshold increased from 10 households to 20. This led to considerably more suppression. Nearly 48 percent of the reported flows between 2000 and 2012 would have been suppressed. For this reason I limit my analysis to 2012. In the analysis sample from 2000 to 2012, 96.6 percent of flows between an origin county and fracking counties in one of the 16 destination states are reported as zero. Some of this likely due to censoring, but it is also probably that many counties do not send any migrants to a distant fracking state in any given year. Among origin/destination pairs that ever report a non-suppressed migration flow, only 39.4 percent of the annual flows are reported as zero. Even among pairs where the smallest non-zero migration flow was less than 20, only 52.5 percent of annual flows are reported as zero. These patterns would be consistent with censoring masking a lot of variation, but with most of the zero reported flows begin true zeros.

Unfortunately, the IRS county to county flows only provide aggregate numbers, and do not break up the migration levels by demographic characteristics (gender, marital status, education). As such, I am unable to use the IRS measures to look at differences across demographics. The only measure provided is the total adjusted gross income for all of the moved- returns. This is the adjusted gross income in the first year, but only the average for all movers in the county pair is provided.

The IRS data does not capture every move from one county to another. Low income individuals and households are not required to file a tax return, and thus might be under represented in the data. It is likely that individuals that move to fracking areas will earn well beyond the filing threshold after moving, but they might not have been required to file in the previous year. If there are individuals that did not file in the first year, but moved in response to fracking and filed in the second year, my estimates would be attenuated. Households that file for extensions past September will also not be included in the data, which might exclude very high income households with complicated returns. Although this is perhaps the most comprehensive data on internal migration in the United States, it might under-represent a subset of the extremely poor (who fall below mandatory tax filing thresholds and do not file for other benefits such as the Earned Income Tax Credit) as well as a small subset of the extremely wealthy (who are more likely to be granted filing extensions for complex returns).

American Community Survey Microdata

To explore heterogeneous responses to news exposure I also exploit the 2005-2012 American Community Survey Microdata obtained from IPUMS (Ruggles et al., 2015). Each year the Census conducts the ACS which surveys approximately one percent of households each year. In addition to collecting information about household structure, demographics (age, race/ethnicity, gender, marital status), education, and employment, it asks individuals where they lived in the previous year. As such, I am able to construct origin/destination migration rates for demographic subgroups. Unfortunately, the geographic data is only available starting in 2005 and the smallest geographic unit is the Public Use Micro Area (PUMA). PUMAs are geographic areas defined by population that are large enough to preserve privacy. Furthermore, migration geographic data is only available at the Migration PUMA (MIGPUMA) level, which are often even larger. These MIGPUMA can contain one or more counties. There are several aspects of the data that are likely to reduce power, making it more difficult to detect an effect. First, because this is based off of a one percent sample of

households there is likely to be measurement error in the constructed migration rates will reduce precision. Second, I have less geographic variation (and less variation in circulation rates) than is available at the county level leading to less precision.

To construct my estimates I use a county to MIGPUMA cross walk and collapse county-level measures to the MIGPUMA level. I then construct MIGPUMA origin to fracking state destination migration flows for each group by calculating the fraction of people in each group from each MIGPUMA origin that moved to a fracking MIGPUMA in a fracking state. This is analogous to the IRS SOI estimation but at the MIGPUMA rather than county level. After Hurricane Katrina, several PUMA and MIGPUMA in New Orleans, Louisiana were combined. I drop individuals living in these combined PUMA in 2005.

Alliance for Audited Media Newspaper Circulation Data

Newspaper circulation rates between 2005 and 2008 were obtained through a temporary academic membership at the Alliance for Audited Media (AAM). These circulation rates are provided in PDFs, which I scraped to collect county level estimates. In some cases the scrap was unable to identify the circulation rate, so hand corrections were made.

The AAM conducts regular (annual or biannual) audits of newspapers and collects circulation rates, along with other measures such as prices. This circulation rates includes the number of copies sold on the audit date and the number of copies as a percent of households for each county with over 25 copies. Counties with fewer than 25 copies sold are assigned a zero value. For most newspapers, circulation rates are reported at the county level. However, for the *New York Times* and *Wall Street Journal* these rates were provided at the DMA-level. For my county level analysis I assign each county the DMA-level, which reduces the variation and adds measurement error. However, as seen in Table 3.5, DMA-level estimates provide similar conclusions. A small subset of local newspapers that reported about fracking do not have AAM audits. For these newspapers, which often only distributed to one or two counties, statistics about local circulation was compiled from online searches. For local newspapers that were not audited annually, the intermittent values were imputed through linear interpolation. The three national newspapers report circulation every year.

TV and Cable Factbook TV Circulation Data

TV circulation data is taken from the Television Cable Factbook for 2008 and 2016. The Factbook contains information on local TV stations as well as DMA-level circulation as reported by Nielsen's. TV circulation is reported at the DMA level for each TV station and includes viewership from both cable and non-cable households. This data is available at the station-level and not specific to news programming. The circulation rate is constructed by dividing total weekly viewership by the total number of households in the DMA. I use average weekly circulation rates throughout my analysis. For each station the "own" DMA and "other" DMA circulation is reported. Because it is not specified what "other" DMA is included, I only include circulation in "own" DMA. This is likely to attenuate the estimated effects. However, for many stations viewership outside the DMA is very low or non-existent. The 2016 circulation rates were obtained through a temporary online membership which provided only the current 2016 circulation rates.

For this reason, I also hired an undergraduate RA to collect circulation rates from the 2008 Factbook. Between 2007 and 2009, TV stations were transitioning from analog to digitally transmitted broadcasts on a market-by-market basis. When a market transitioned, viewers were required to obtain digital reception equipment, and it is unclear how this affected viewership and if 2008 viewership is representative of later years. For this reason I include estimates using both the 2008 and 2016 measures.

LexisNexis Newspaper Content Data

Newspaper content is collected through LexisNexis by searching on key terms, “frack*”, “fracing”, and “hydraulic fractur*”. I then take the universe of articles, remove non-US sources (e.g., Daily Mail in the UK), and remove articles that only reference things like “Frick and Frack”, unrelated acronyms, or last names. I then parse the entire text of these articles for each of the 16 state names (both capitalized and lower cased). References to states in the title of newspapers or place of publication are excluded, (ex: articles published in Colorado are not included as citing Colorado unless there is a reference in the body of the text). I then parse the entire text of the articles for positive and negative terms: “new job”, “creat + job”, “low + unemploy”, “hire”, “hiring”, “boom”, “growth”, “earthquak”, “environment”, “health”, “contaminat”, “danger”, and “pollut”. Positive articles are articles that reference at least two positive terms and more positive terms than negative. Negative terms are the opposite. There are “neutral” articles that refer to fewer references that are not included when looking at news content, but have been included in previous specifications. When positive, neutral, and negative news are all included the patterns are similar but less precise.

Vanderbilt Television News Archive TV News Content

TV news content was pulled from the Vanderbilt Televisions News Archive (VTNA) and includes broadcasts that mention “fracing”, “frack”, or “shale”. The VTNA database contains TV news recordings and transcript abstracts for nightly news broadcasts from the three major news networks (ABC, CBS, and NBC) and the cable news channels CNN and Fox News. The database only includes one hour of programming each day for both cable news outlets. Because the available content of cable news is limited, and circulation rates are only available for the TV networks, I restrict the sample to TV broadcasts from the three major news networks. I parse the transcript abstracts for search terms such as “fracking” and “shale” as well as which state is being discussed. These clips are short often ranging from one to five minutes in length.

LEHD Origin Destination Employment Statistics Commute Data

The LEHD Origin Destination Employment Statistics (LODES) contains the number of workers for each residence/work place Census block pair. This data is available for all years since 2002, and also provides statistics by broad age (under 30, 30-54, over 54), monthly earnings (under \$1,250, \$1,250-3,333, over \$3,333), and industry (goods, trade/transportation, other) groups. For each Census block I identify the corresponding county, and then aggregate up commute flows to the county to county level. For privacy, some noise is introduced at the Census block level, which likely remains at the county level, although to a lesser extent.

Google Trends Data

For a given search term, Google will provide a measure of search intensity within a given geographic area over a specified time period. The geographic regions include the entire country, states, or the DMA-level. All of my analysis is conducted using DMA-level search indices. For a given term (like “fracking”) Google will identify the day in the search window where the term is the most searched as a fraction of all searches. This day is assigned a value of 100. All other days will be indexed to that day. So a value of 20 would represent a search intensity of 20 percent. Because these indices are produced within the search area and time they are internally consistent but not comparable across queries. If there is not enough search for a given term over the period, no values are reported. I have also looked at search interest in moving specific terms such as

“Uhaul” or “Uhaul rental”. At both the DMA and state-level there appears to be a visual shift at the time of broadcast, however it is not statistically different. Search intensity for terms like “fracking jobs” or “oil jobs” are low and frequently suppressed by Google. These patterns are similar if the window is extended.

Twitter Tweet Data

Twitter provides an API to access tweets, but only recent tweets (in the last 14 days) are recoverable. Because I am interested in historic tweets, I scrape them directly from Twitter using the `GetOldTweets3` package in Python. Using this method I can only recover location if users include this with their post, and in practice this happens rarely. Instead, I consider the total tweet count for the entire country at the daily level. Using the `GetOldTweets3` package, I collect all tweets that include the word “fracking” anywhere in the text. I also observe the entire text, the day and time, the number of likes and re-tweets, the sender, and whether any other Twitter users handles (@) were specifically mentioned. Tweets that are re-tweets (without additional text) are not included as separate tweets and only counted as re-tweets.

For Online Publication: Appendix C. Conceptual Model: Information and Migration

In the canonical migration choice model (Sjaastad, 1962), an individual will move if the lifetime utility derived from moving minus the fixed costs of moving exceeds the utility of staying at the original location. The individual observes the real returns $y_d(t)$ and $y_o(t)$ for each period in each location (which account for earnings, cost of living, local amenities, and idiosyncratic fit) as well as the fixed utility cost c_{od} associated with migrating from o to d . These location specific returns can vary over time, and are discounted by β . Under complete information with no uncertainty, the decision to move from o to d (m_{od}) is as follows

$$m_{od} = \begin{cases} 1 & \text{if } \sum_{t=0}^T \beta^t u(y_d(t)) - c_{od} \geq \sum_{t=0}^T \beta^t u(y_o(t)) \\ 0 & \text{else} \end{cases} \quad (10)$$

However, individuals likely face incomplete information about conditions in the potential destination. As such, the individual views $y_d(t)$ as a random variable, where $y_d(t) \sim G(y; \theta)$.⁴⁹ Now, the individual will only migrate if

$$\sum_{t=0}^T \beta^t (Eu(y_d(t)) - u(y_o(t))) - c_{od} \geq 0 \quad (11)$$

where the E represents the expected value at time zero. Equation (2) yields a threshold moving cost (c_{od}^*) at which the individual is indifferent between staying and moving

$$c_{od}^* = \sum_{t=0}^T \beta^t (Eu(y_d(t)) - u(y_o(t))). \quad (12)$$

Changes in the parameters θ which govern the distribution of $y_d(t)$ will affect this cost threshold. For simplicity, consider the case where $y_d(t)$ is distributed normally with a mean (μ_d) and variance (σ^2) such that $\theta = \{\mu_d, \sigma^2\}$, and that the individual is risk averse with monotonic preferences (e.i., $u' > 0$ and $u'' < 0$). Under these assumptions

$$\frac{\partial c_{od}^*}{\partial \mu_d} = \sum_{t=0}^T \beta^t \frac{\partial Eu(y_d(t))}{\partial \mu_d} > 0 \text{ and } \frac{\partial c_{od}^*}{\partial \sigma^2} = \sum_{t=0}^T \beta^t \frac{\partial Eu(y_d(t))}{\partial \sigma^2} \leq 0. \quad (13)$$

Intuitively, as the mean increases, the individual places less weight on low values of y_d leading to higher expected utility and the individual is now willing to pay a larger moving cost. As Rothschild & Stiglitz (1970) show, an increase in the variance, holding all else equal, represents a mean preserving spread which results in weakly lower expected utility because the individual is risk averse.⁵⁰ The increase in variance leads to more risk, and the individual's moving cost threshold weakly falls, as she must be compensated by a lower cost to move to compensate for the added risk. If the individual's prior belief is that the return to migration

⁴⁹The model implications are similar if the individual lacks information about conditions at the origin.

⁵⁰If the distribution of $y_d(t)$ is governed by more than just locational parameters this is not necessarily true (Tobin, 1965; Dionne & Harrington, 1991). More generally, if $\hat{\sigma}^2$ is a mean preserving spread of σ^2 , then Rothschild and Stiglitz (1970) prove that $Eu(y_d(t); \sigma^2) \geq Eu(y_d(t); \hat{\sigma}^2)$. If instead the utility is linear and individual is risk neutral, changes in the dispersion that preserve the mean will not affect the cost threshold.

is low, she will be less willing to move (as seen in equation (13)). Similarly, if her prior is diffuse and the investment in migration appears more risky, she will also be less willing to move.

People might have incomplete information about the parameters that govern the distribution of $y_d(t)$ (μ_d and σ^2).⁵¹ Receiving additional pieces of information about these parameters can lead individuals to update their beliefs about μ_d and σ^2 , and potentially change migration behavior. This seems like a plausible scenario in the case of fracking where there was often little known about the places involved in fracking and there was a lot of information about fracking provided in the news. This corresponds to a framework of updating, where people recognize that they are missing information and have prior beliefs about the distribution of μ_d and σ^2 that they update upon receiving new information. This updating process will impact their perceived distribution of $y_d(t)$.

Suppose $y_d(t) = \mu_d + \varepsilon_d$, where μ_d and ε_d are independent and $\varepsilon_d \sim N(0, \sigma^2)$, but the decision-making individual has incomplete information about the mean of $y_d(t)$.⁵² Given her prior observations of the world, she believes $\mu_d \sim N(\mu_{d0}, \lambda_0)$. Now suppose she is exposed to information (perhaps in the newspaper or through TV news) that gives information about $G_0 = (g_1, g_2, \dots, g_{n_0})$, draws from the distribution of $y_d(t)$, which are assumed to be independent. The individual can use this information to update their beliefs about the distribution of μ_d using a process like Bayesian Updating. According to Bayes Rule

$$P(\mu_d|G_0) = \frac{P(G_0|\mu_d)P(\mu_d)}{P(G_0)} \quad (14)$$

Because $P(G_0)$ is fixed constant, it is also true that

$$P(\mu_d|G_0) \propto P(G_0|\mu_d)P(\mu_d). \quad (15)$$

In other words, the individual's perceived distribution of μ_d , conditional on observing G_0 is proportional to the likelihood of observing G_0 given the parameter μ_d times the prior perceived distribution of μ_d . By observing both the likelihood of G_0 and the prior of μ_d , the individual can update her perception of the distribution of μ_d , which is known as the posterior. If she received new information sets (G_0, G_1, G_2, \dots) over time, this updating process can be repeated iteratively, allowing the individual to incorporate the new information and adjust her beliefs about the distribution of μ_d . For certain distributions, such as the normal distribution, the posterior and prior probability distributions from Bayesian Updating are in the same family of distributions, and the parameters of the posterior distribution are formulaically adjusted by the new data observed. In this case, the posterior distribution will be

$$P(\mu_d|G_0, \sigma^2) = N\left(\mu_d \left| \frac{\sigma^2}{n_0\lambda_0 + \sigma^2}\mu_{d0} + \frac{n_0\lambda_0}{n_0\lambda_0 + \sigma^2}\bar{x}, \frac{\sigma^2\lambda_0}{n_0\lambda_0 + \sigma^2} \right| \right) \quad (16)$$

where \bar{x} is the sample average, n_0 is the number of data points, μ_{d0} is the mean from the prior, σ^2 is the true variance of $y_d(t)$ and λ_0 is the variance from the prior probability distribution.⁵³ From this updating

⁵¹This type of incomplete information is prevalent. Even among highly educated medical students in the residency match process there is substantial heterogeneity in their ability to accurately predict the expected cost of living and earnings rank in their top two ranked locations (Bottan & Perez-Truglia, 2017).

⁵²In this scenario, the individual is assumed to know the variance of $y_d(t)$. A similar process arises if both the mean and variance are unknown, but now the posterior will vary across the two parameters.

⁵³This is a general result in the Bayesian Updating literature (Murphy, 2007).

process, the individual can identify a new updated posterior mean and variance for the unknown distribution of parameter μ_d

$$\mu_{d1} = \frac{\sigma^2}{n_0\lambda_0 + \sigma^2}\mu_{d0} + \frac{n_0\lambda_0}{n_0\lambda_0 + \sigma^2}\bar{x} \quad (17)$$

$$\lambda_1 = \frac{\sigma^2\lambda_0}{n_0\lambda_0 + \sigma^2} \quad (18)$$

Although the individual is likely not performing these calculations to incorporate new information, using sample moments from the new information to update beliefs is reasonable at an intuitive level.⁵⁴ The posterior mean is a weighted average of the prior mean and the sample mean from the new draw of data. The weights correspond to the relative dispersion. An increase in the sample mean from the new information will lead to a larger posterior mean. The magnitude of this increase will depend on how precise or diffuse the individual's prior belief is. If the individual has a diffuse prior (λ_0 is small), she will attribute the difference between μ_{d0} and \bar{x} to observational noise and not change her beliefs in response to the new information. If the individual is uncertain about her prior (λ_0 is large) she will put more weight on the new data and take differences in the sample average as a signal that her prior needs to be adjusted.

Notice that λ_1 , the posterior variance of μ_d , is decreasing in n_0 . In other words, as more pieces of information and data are received the individual becomes more confident in her prior. This will lead them to put less weight on future data as they become more confident in their prior. As such, the marginal impact of each additional piece of data will become smaller as more data is received, potentially leading to non-linear, decreasing returns to new information. The impact of each piece of information will depend on her confidence in her initial belief about μ_d .

As the individual updates her beliefs about μ_d this will shape her belief about $y_d(t)$, the return to moving to destination d . As μ_d changes her belief about the distribution of $y_d(t)$ can be computed by determining the likelihood of observing some return y given data points G_0 as

$$\begin{aligned} P(y|G_0) &= \int P(y|\mu_d)P(\mu_d|G_0)d\mu_d \\ &= \int N(y|m_d, \sigma^2)N(\mu_d|\mu_{d1}, \lambda_1)d\mu_d \\ &= N(y|\mu_{d1}, \lambda_1 + \sigma^2). \end{aligned} \quad (19)$$

We can now determine how information affects the distribution of $y_d(t)$ (see Bishop (2006) for proof). The effect of additional information on the perceived mean of $y_d(t)$ will interact with prior beliefs. If the prior belief about average returns in a potential destination was lower than the news suggests, (a likely case in this scenario given that many fracking locations were rural and unknown and that much of the coverage of fracking was positive), additional data will increase μ_{d1} . As we see in equation (13) this will result in a higher moving cost threshold as the individual becomes more willing to move. The opposite case is also possible. If the individual's prior belief was that the average return in fracking locations were higher than portrayed in the news, new information will result in a lower μ_{d1} in the updating process, making migration less desirable.

⁵⁴Wiswall and Zafar (2015) show that when college students receive information about the distribution of earnings, they update their beliefs, but often do not strictly follow a Bayesian updating process.

Obtaining more data will always reduce the dispersion of $y_d(t)$ by reducing λ_1 (confidence in prior of the mean). However, there will still be uncertainty associated with the distribution of returns (σ^2). Shifting upward her prior about the mean, holding all else equal makes moving less risky and will make migration more desirable for risk averse individuals as in equation (13).

People’s responses to news about localized fracking booms might have followed a similar process. Exposure to news stories that credit fracking with creating local booms, fueling local economic growth, or raising wages in potential destinations might change people’s perceptions of the distribution of the returns in the fracking destinations mentioned; even negative news can provide information about where fracking is occurring and change people’s beliefs.⁵⁵ For example, individuals exposed to numerous newspaper articles and TV news broadcasts touting the local economic benefits of fracking in Texas might adjust their beliefs about the average or dispersion of the return to migrating to a Texas fracking county. This news information does not necessarily need to be correct, as long as the individual believes it contains truthful information.

As the individual incorporates new information about the parameters (μ_d or σ^2), she better understands the distribution of $y_d(t)$ and can compute the likelihood of observing the return y given her information set. The effect of additional information on the perceived mean of $y_d(t)$ will depend on her prior beliefs. If she initially believed the average return in a potential destination was lower than the news suggested (a likely case given the coverage about fracking jobs and booms), the information will increase her perception of μ_d . This in turn increases c_{od}^* , meaning she is more willing to move (see equation (13)).

Simulated Effects of News Information on Updating

To visualize how information can change migration decisions when individuals lack information a simple simulation is shown in Appendix Figure C1. Two scenarios are presented for three types of people. Individual 1 has a diffuse prior over the expected return (μ_d) and incorrectly believes μ_d is lower than the true mean. Individual 2 has a precise belief that μ_d is low. Individual 3 has a diffuse prior, but correctly predicts μ_d . In both scenarios the true parameters are the same, the only difference is that individual are exposed to more information in scenario 1 than in scenario 2.

The perceived distribution of both μ_d and y_d are plotted for each individual in each scenario over two iterations of receiving more “news”. If initial beliefs about the expected return are low, new information shifts up the beliefs about μ_d and y_d . Additional information also reduces dispersion of μ_d and y_d , which increases expected utility and the probability of migrating. Updating is more drastic when there is more information, and changes in the probability of moving will be the largest among people or areas that are exposed to more new information. Initial draws of information are very beneficial, but the marginal value of additional information becomes smaller.

⁵⁵Up through 2012, the last year of my sample, about 60 percent of adults were familiar with fracking, and over half of this population was in favor of fracking (Pew Research, 2013a). For someone that views fracking favorably, even a negative news story could provide information about where fracking is occurring, and result in updated beliefs.

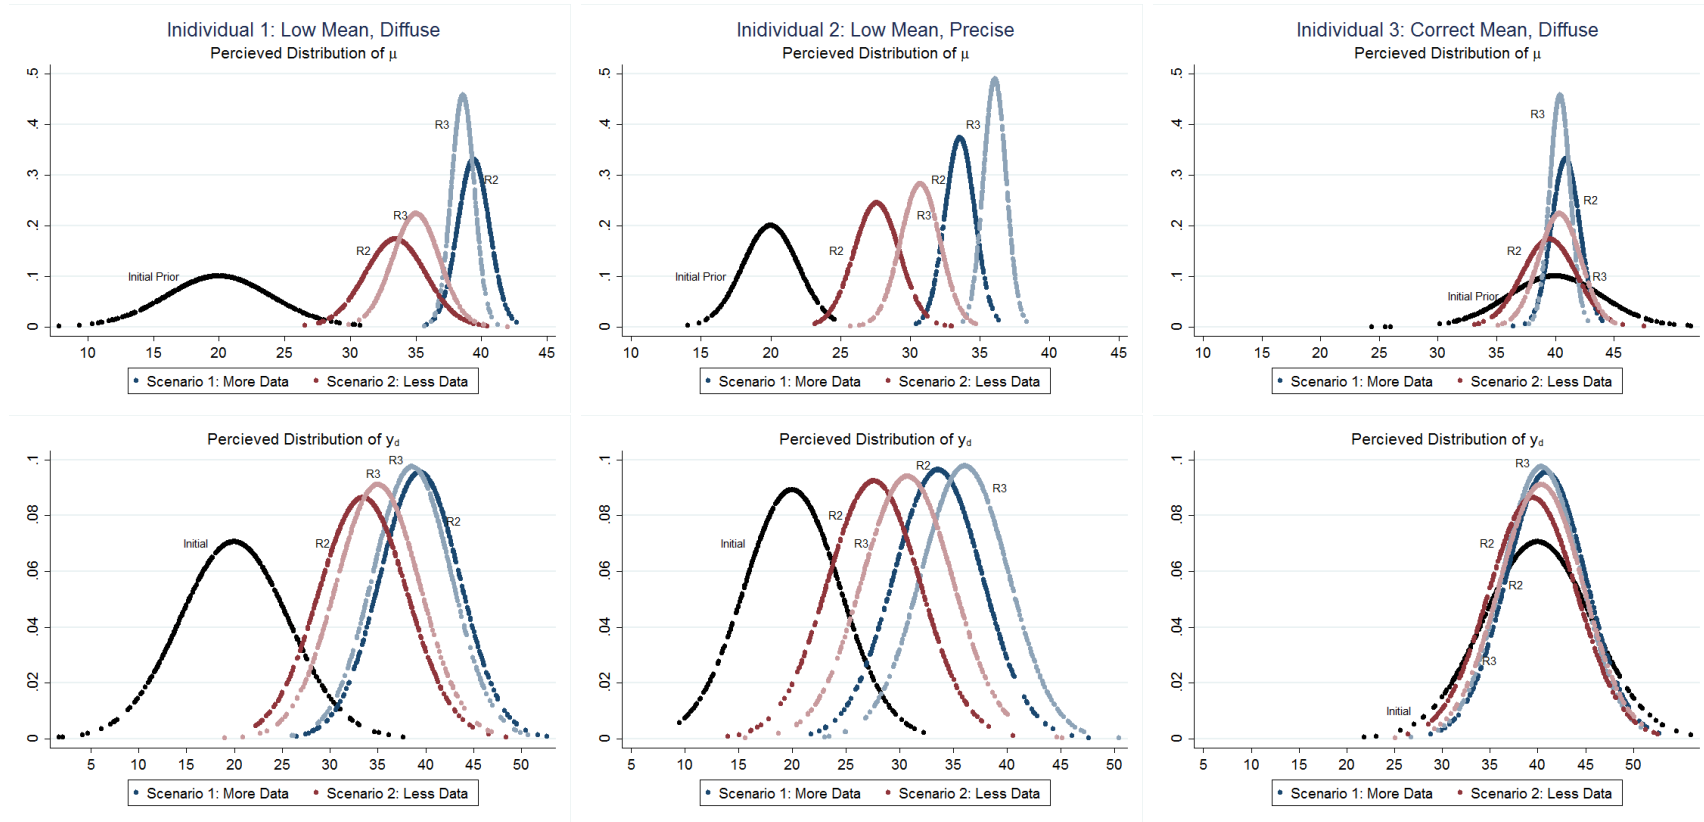


Figure C1

Model Simulations: Information and Bayesian Updating

Notes: Simulated data points from the distributions of μ_d and y_d are presented for three separate individuals in two separate scenarios. Individual 1 had a diffuse prior with a low mean, individual 2 had a more precise prior with a low mean, and individual 3 had a diffuse prior with a correct mean. In scenario 1 the individual viewed ten data points from the true distribution of y_d in each round (R2 and R3), and updates the posterior probability accordingly. In scenario 2 the individual views only 2 data points and updates the posterior. The initial prior and two additional iterations are shown.

Source: Author's calculations.