Moving Back Home: Insurance against Labor Market Risk

Greg Kaplan

University of Pennsylvania

This paper demonstrates that the option to move in and out of the parental home is a valuable insurance channel against labor market risk, which facilitates the pursuit of jobs with the potential for high earnings growth. Using monthly panel data, I document an empirical relationship among coresidence, individual labor market events, and subsequent earnings growth. I estimate the parameters of a dynamic game between youths and parents to show that the option to live at home can account for features of aggregate data for low-skilled young workers: small consumption responses to shocks, high labor elasticities, and low savings rates.

I. Introduction

Typical life cycle models, which are used to study how individuals cope with labor market shocks, abstract from the possibility that young adults can live with their parents. Yet many youths continue to live with their parents after entering the labor market and often move back in with

I thank Gianluca Violante, Thomas Sargent, and Chris Flinn for the valuable advice they have provided at each stage of this project. This paper is based on chap. 3 of my PhD dissertation at New York University. I am grateful for the comments and suggestions from Ellen McGratten, Monika Piazzesi, and three anonymous referees. I have benefited from participants at seminars and from discussions with people at many institutions that are too numerous to mention. Part of this research was conducted while at the Federal Reserve Bank of Minneapolis. Any views expressed here are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis of the Federal Reserve System. Diego Amador provided excellent research assistance.

[*Journal of Political Economy,* 2012, vol. 120, no. 3]
© 2012 by The University of Chicago. All rights reserved. 0022-3808/2012/12003-0003$10.00

their parents after having lived away from home for a period of time.[1] In this paper, I use an estimated structural model to account for these coresidence dynamics and show that the option to move back home is a valuable channel of insurance against labor market risk. The use of coresidence as insurance has important quantitative implications compared with settings from which the possibility of parental coresidence is abstracted: smaller consumption responses to job loss, higher labor elasticities of the young, lower savings rates, and higher long-term earnings growth.

In my model, parental coresidence is determined by the outcome of a dynamic game between youths and parents. Parents, who are altruistic toward their offspring, can provide both monetary support, through explicit financial transfers, and nonpecuniary support in the form of shared residence. The benefits of shared residence accrue from a reduction in per capita direct housing costs and the availability of public goods inside the parental home. There are psychic costs of shared residence due to the lack of independence. Youths make labor supply and savings decisions, in addition to choosing whether to live with their parents. To allow for the possibility that there are long-term labor market consequences of the option to live at home, the model features two types of jobs: one that is low risk and easy to find but generates low earnings growth and another that is risky and harder to find but is associated with better earnings growth possibilities. The allocations of interest are given by the unique Markov-perfect equilibrium (MPE) of this game.

To simultaneously account for the cross-section and time series dimensions of coresidence and labor market outcomes in the data, the model requires two types of idiosyncratic uncertainty. The first are labor market shocks that come in the form of stochastic job offers, job losses, promotions, and changes in productivity. These shocks generate a motive for moving in and out of home since the housing decision and the labor market are tied together through a budget constraint. The second are shocks to youths' relative desire to live away from their parents. These reflect non–labor market events: social factors, finding a partner, or maturity. The quantitative implications of parental coresidence for youths' behavior—either through short-term insurance against shocks or by facilitating longer-run earnings growth—depend crucially on whether observed coresidence dynamics are mostly due to the labor market shocks or the preference shocks.

[1] A substantial body of anecdotal evidence and reports in the popular press suggest a recent trend in the United States for young people to move back home with their parents after a period of living away from home. This has led to the coining of the term "Boomerang Generation" to describe this group. See the Wikipedia entry at http://en.wikipedia.org/wiki/Boomerang_Generation and the numerous references cited there.

I structurally estimate the model using monthly data for young males who do not go to college from the National Longitudinal Survey of Youth 1997 (NLSY97) and find that labor market shocks affect the timing of when youths move in and out of their parents' homes, while preference shocks affect cross-sectional differences in living arrangements. The high-frequency nature of the NLSY97 panel data is crucial in identifying the impact of the two shocks since it allows me to observe labor market outcomes and events around the time that coresidence transitions take place. Specifically, I find that in any given month the difference in earnings between youths living at home and away is small but that for each individual youth, the timing of when they move in and out of their parents' home is closely related to their labor market experience. Together these facts yield an estimate for the unobserved process for preference shocks that is very persistent but with a large amount of cross-sectional heterogeneity. This implies that the extent to which labor market shocks account for coresidence patterns differs for the cross section of living arrangements and the within-individual time series of living arrangements: whereas only a small fraction of the cross-sectional differences are accounted for by labor market outcomes, the majority of movements in or out of the parental home are driven by labor market events. This is an important distinction—it implies that the importance of parental coresidence as an insurance channel would be overlooked if one were to restrict attention to cross-sectional regressions or a static structural model.[2]

I use the estimated model to measure the value of different insurance channels, by comparing the welfare cost of a job loss with the corresponding welfare cost when a particular insurance channel is removed. I find that the option to move in and out of home is valuable for all youths, but particularly so for youths from poor families: the welfare cost of a job loss for a youth in the bottom quartile of the parental income distribution is 12 to 20 times larger without the option to move back home, depending on the type of job that is lost, and is three to four times larger for a youth in the top quartile. This is because parents from the lower part of the income distribution find it more costly to

[2] The existing empirical literature has largely restricted attention to cross-sectional patterns of coresidence at a point in time, or the first movement out of the parental home, focusing on comparing the family and individual characteristics of youths living at home versus youths living away from home. Prominent examples include McElroy (1985), Buck and Scott (1993), Rosenzweig and Wolpin (1993, 1994), Card and Lemieux (1997), Ermisch and Di Salvo (1997), and Manacorda and Moretti (2006). There are very few empirical analyses of movements back home. Three exceptions are Da Vanzo and Goldscheider (1990), who use annual data from the 1972 National Longitudinal Surveys; Goldscheider and Goldscheider (1999), who use retrospective information on whether an individual ever moved back home from the National Survey of Families and Households; and Ermisch (1999), who uses annual data from the British Household Panel Survey.

substitute financial transfers for coresidence when the option to move back home is removed, whereas for youths with wealthier parents, financial transfers are a close substitute for coresidence.

Through a set of counterfactual exercises, I then show that the option to move back home has important quantitative implications for consumption, labor supply, and savings behavior. In the estimated model, the consumption drop in response to a job loss is around 47 percent lower than it is when youths are precluded from living with their parents. This suggests that incorporating parental coresidence could improve the fit of existing models in terms of how much consumption responds to labor market shocks. For example, Kaplan and Violante (2010) compute the amount of consumption insurance implicit in a calibrated incomplete-markets life cycle economy without coresidence, where self-insurance through borrowing and savings is the only private insurance channel, and compare it with corresponding estimates from US data in Blundell, Pistaferri, and Preston (2008). They find that it is young households with low wealth for whom the hypothesis of self-insurance alone is most at odds with the data. This group of households is particularly relevant when considering coresidence as an insurance margin.

The estimated model is consistent with a high aggregate labor elasticity of young workers, thus addressing a recent literature that notes that labor market fluctuations at business cycle frequencies are disproportionately large for this group compared with older workers (Rios-Rull 1996; Gomme et al. 2004; Jaimovich and Siu 2009). Recognizing that for many young people, the outside option when making their labor market decisions includes the opportunity to live with their parents leads them to raise their reservation wages to a point where job acceptance rates are more sensitive to changes in the overall distribution of wage offers than if youths were precluded from living at home. Without the possibility of parental coresidence, job acceptance probabilities are around 15 percent higher in the model.

I find that when the option to move back home is removed, asset accumulation increases by around 16 percent. Hubbard, Skinner, and Zeldes (1995) show that asset-based, means-tested social insurance, such as that implicit in the US welfare system, can have distortionary effects on savings behavior, by discouraging households from accumulating their own precautionary wealth. In my model, a similar effect operates for young males through the implicit insurance provided by their parents. This highlights the fact that the option to move back home can have important behavioral implications for youths living at home or away, even if they never actually experience such a move. Simply knowing that this opportunity exists reduces the precautionary incentives to accept jobs or accumulate financial assets.

Finally, the option to live at home has a substantial impact on youths' future earnings, by providing the means to search for jobs with higher potential for earnings growth but that are risky, take longer to find, or provide lower earnings in the short term. When the possibility of coresidence is restricted, youths are more likely to accept easy-to-find safe jobs with low earnings growth, and hence average earnings growth is lower. In the model, average monthly earnings at age 23 are around 5 percent higher when youths have the possibility of living with their parents than when they do not.

The prediction that the option to move back home leads to higher future earnings is consistent with new empirical evidence from the NLSY97 that I document independently of the model. By comparing labor market outcomes at age 26 with coresidence and employment dynamics at age 20, I show that a job separation at age 20 is associated with a large reduction in earnings 6 years later. However, this reduction is only present for youths who were not living with their parents at the time of the job separation and is greatly reduced for youths who were living away from their parents at the time of the job separation but moved back home in the subsequent 3 months.

Another empirical contribution of the paper is to show that the relationship between parent-child living arrangements and labor market outcomes also extends to older individuals of both genders in all education groups. Using data from the Current Population Survey (CPS), I show that the aggregate parental coresidence rate moves closely with aggregate labor market conditions, for all individuals up to age 34. To provide evidence for causality, I then use differences-in-differences across US states to show that a decline in state-level employment or hours worked leads to a significant increase in the state-level parental coresidence rate. Thus, although the structural estimation and the empirical analysis with the NLSY97 data focus only on young males aged 17–22 who do not go to college, the quantitative importance of parental coresidence likely extends to a much larger part of the population.

There is a related empirical literature that has also studied the determinants of living arrangements. This literature is best exemplified by Rosenzweig and Wolpin (1993, 1994), who use annual data to estimate logit and multinomial logit models with fixed effects for the probability of youths residing with their parents, receiving financial transfers from their parents, and receiving transfers from the government. On the basis of their empirical findings, they argue that parental coresidence is a commonly used form of intergenerational support. While these papers are useful for understanding the determinants of coresidence in the cross section, my analysis with monthly panel data allows for the estimation of duration models with unobserved individual-specific effects that relate the hazard of moving in and out of home to the labor market. Such an empirical approach is more useful for understanding the effects of the labor market on coresidence transitions. The existing literature does not provide the necessary empirical evidence, or the quantitative dynamic model, that is needed to measure the implications of coresidence for youths' behavior that I do here. Furthermore, in this paper the empirical analysis of coresidence transitions with high-frequency data, together with the estimated structural model, permits the construction of counterfactuals and makes it possible to study the impact of coresidence over both the short and the long term.[3]

II. Evidence on Parental Coresidence and the Labor Market

I start by providing empirical evidence on the relationship between parental coresidence and labor market outcomes using two types of data. First, I use monthly individual-level panel data to establish a direct relationship between labor market events and movements in and out of the parental home. I show that coresidence transitions are common and are strongly influenced by employment and earnings outcomes, suggesting a possible role as a mechanism for insuring the short-term effects of labor market shocks. I then show that the option to move in and out of home may also have important long-term implications for labor market outcomes, by helping to mitigate the effects of job displacement on earnings. Second, I use a difference-in-difference approach to provide evidence that causality does indeed run in the direction of labor market shocks to coresidence transitions. I exploit state-level variation in aggregate hours and employment using repeated cross-section data from the CPS to show that when labor market conditions deteriorate, parental coresidence rates increase significantly. I

[3] The idea that families have an important role to play in smoothing the impact of economic shocks dates back at least as far as the seminal work of Becker (1974). For an excellent review of work on the various forms of intergenerational ties in economics, sociology, and psychology, see Bianchi et al. (2006). The work in this paper is related to this and a number of other branches of existing literature: (i) studies of the provision of support after retirement and into old age, including Costa (1999), Pezzin, Pollak, and Schone (2007), and Bethencourt and Rios-Rull (2009); (ii) the purely theoretical analyses of youth coresidence in Fogli (2004) and Becker et al. (2010), which allow for expectations about future outcomes to affect current coresidence decisions; (iii) the class of models that incorporate a savings decision into a labor market search setting as in Danforth (1979), Lentz and Tranaes (2005), Low, Meghir, and Pistaferri (2010), and Lise (2011); (iv) the line of literature originally advocated by Deaton and Paxson (1994) that attempts to quantify the extent to which idiosyncratic shocks are insurable (Attanasio and Davis 1996; Hayashi, Altonji, and Kotlikoff 1996; Storesletten, Telmer, and Yaron 2004; Blundell et al. 2008; Kaplan and Violante 2010); (v) the literature that studies the effects of "real-world" channels on consumption, which includes Fernandez-Villaverde and Krueger (2002) (durable consumption goods), Low (2005) and Kaplan (2011) (variable labor supply), and Livshits, MacGee, and Tertilt (2007) (bankruptcy protection). Sakudo (2007) also estimates a structural model of coresidence for young females in Japan, focusing on the decision to move out of home and get married.

also show that this effect is evident at a national level during the large recessions of 1982 and 2008.

A. *Panel Data Evidence from NLSY97*

Panel Data for Measuring Coresidence

The NLSY97 is a longitudinal survey of 8,984 individuals from the cohort born between 1980 and 1984. They have been sampled approximately annually since 1997. The survey contains extensive information on each youth's labor market behavior and educational outcomes, together with detailed information about family and community background. What makes the NLSY97 an ideal data set for studying the dynamics of parent-youth living arrangements is a set of retrospective questions about monthly coresidence that were asked in rounds 2–6 (1998–2002). At each interview, these questions asked respondents to list each period of 1 month or more in which they lived separately from each of their parents, where a parent is defined as a biological, step, adoptive, or foster parent. From these questions, it is possible to reconstruct a monthly panel of parental coresidence outcomes for each respondent, which I then merged with data on education, labor market, and marital histories.[4] The resulting data set is the first to contain high-frequency (greater than annual) data on parental coresidence and the only one that contains information on labor market outcomes at the times that coresidence transitions take place. It thus provides a unique opportunity to understand the circumstances that surround changes in youths' living arrangements.

Sample Selection

Sample selection poses a challenge for this analysis. First, the fact that the monthly coresidence questions were discontinued in 2002 restricts the ages at which it is possible to observe contemporaneous labor market and coresidence outcomes in the NLSY97. In particular, this means that studying the interaction between labor market dynamics and coresidence dynamics for youths who go to college is not possible. For my

[4] Since the monthly information on parental coresidence is retrospective, there is the potential for recall bias to affect the results. I do not find any systematic patterns in coresidence states or transitions as a function of the recall time, suggesting that recall bias is unlikely to be large enough to significantly affect the analysis.

analysis with the NLSY97, I thus focus attention on the population of low-skilled youths who do not attend college.[5]

However, implementing the restriction to the noncollege population raises its own challenges. First, for youths in this age group, the decision about whether to attend postsecondary education is likely endogenous with respect to labor market opportunities and coresidence outcomes. Hence, selecting on the basis of observed education choices may introduce nonrandom selection on unobserved characteristics into the sample. Second, many youths may initially decide to enter the labor market in the years immediately after high school, but they may return to education at some point in the future. Moreover, it is common for youths to attend nontraditional part-time colleges, a decision that may also be correlated with labor market opportunities.

My approach is to choose a baseline sample of youths who are never observed to participate in any type of postsecondary education. Choosing a sample of youths for whom we can condition on the decision to not attend college allows the focus to be placed clearly on the interaction between residential movements and labor market events. It seems a natural starting point for understanding the economic implications of coresidence movements for low-skilled youths and avoids the complications that arise from the interaction with college choice.

To address the concerns regarding potential endogeneity of the education decision, I also compare the baseline sample with two alternative samples that implement the restriction to low-skilled youths in different ways. First, I select on the basis of low test scores, which are a strong predictor of future college participation. Since selection into this sample is based on a purely exogenous variable, there are no issues of endogeneity of education. Second, I construct a less restrictive sample by only dropping youths who are traditional college participants—those youths who start college immediately after graduating from high school or within 1 year of graduating—thus retaining youths who may attend college part-time in the sample. I find that both of these alternative samples generate statistics that are very similar to the baseline sample (see table A2), indicating that none of the structural estimation results would be severely affected by selecting low-skilled youths using a different criterion.[6]

[5] After 2002 (round 6), the retrospective coresidence questions were replaced with two questions that ask about the month and year that a youth first lived away from his or her parents and the month and year when he or she returned home for at least 3 months. It is possible, as the cohort ages, that these questions could also be used to study movements back home for college graduates.

[6] This finding does not imply that education decisions are exogenous with respect to coresidence choices. The quantitative conclusions are largely unaffected because, conditional on not being in education, the interaction between coresidence and the labor market turns out not to be affected by previous education decisions.

The other important selection criteria are as follows. Females are dropped, as are males who ever go to the military or have all parents dead.[7] A youth is included in the final panel from the first month after he stops attending high school or after he turns 17, whichever is later. Only youths who have nonmissing residence data are included in the sample. The final sample consists of 36,222 month-youth observations, for 1,491 male youths ranging in age from 17 to 23. These generate 351 spells back home, where a spell is defined as one that is not left censored. Since the NLSY97 has an oversampling of black and Hispanic youths, sample weights are used in all calculations. Tables A1 and A2 report the number of respondents lost at each stage of the selection process and the statistics from the three samples that form the basis for the structural estimation in Section IV, respectively.

Moving Back Home

It is common for young males to move back home after initially having left home: 40 percent of 22-year-olds in the sample have moved back home for 1 month or more. To move back home, one must first live away from home, and 68 percent of youths in this age group are observed to have lived away from home at some point. This implies that 58 percent of the youths who were at risk of moving back home do actually return home. On average, those youths who do move back tend to stay home for a substantial period: the median duration back home is 16 months. However, there is considerable heterogeneity in spell durations. Some spells are fairly short (22 percent end within the first 6 months), yet others are very long (34 percent are 2 years or more).

In the NLSY97, it is difficult to measure the distance that youths move when they move back home since geographic information is only collected on annual interview dates and only for the residence that the youth considers to be his permanent home. However, by focusing on youths who moved only once between interview dates, we can infer that for many youths these moves involve large distances. For this subsample, 41 percent (16 percent) of moves out of home, and 38 percent (18 percent) of moves back home, involved a move to a different city (state).

[7] Focusing exclusively on males avoids modeling complications due to benefit eligibility for young females. Whereas very few males of this age group receive some form of government support, it is not uncommon for females to receive benefits. Many such benefits are related to early childbearing, and, importantly, eligibility for certain benefits is conditioned on residential choices. While this is an interesting feature of the benefit system for females that could be explored in future research, I focus on the simpler case of males in this paper.

TABLE 1
DISCRETE CHOICE MODELS FOR LIVING AWAY FROM PARENTS
(Dependent Variable: Indicator for Whether Youth Lives Away from Parent)

	LOGIT				FIXED-EFFECTS LOGIT	
	(1)	(2)	(3)	(4)	(5)	(6)
Working	1.236**	1.302***			1.100	
	(.119)	(.122)			(.153)	
Log earnings			1.077**	1.115***		1.082
			(.039)	(.035)		(.069)
Away $(t-1)$		908.966***		951.515***		
		(76.903)		(94.420)		
Includes fixed effects					Yes	Yes
Observations	36,222	34,731	26,034	25,125	17,412	12,315
Individuals	1,491	1,479	1,364	1,353	531	458

NOTE.—All models include controls for race, high school graduation, marriage, biological children, parental education, parental income, biological parents married, cubic polynomial in youth age, quadratic polynomial in average parent age, and region dummies. Standard errors are in parentheses, clustered by individual. Parameters are multiplicative effects of probability of working, or marginal change in earnings, on probability of living away from parents.

** Significant at 5 percent level.

*** Significant at 1 percent level.

Discrete Choice Analysis

I begin my analysis of the relationship between coresidence and the labor market with a simple discrete choice logit model for living away from home. Table 1 reports logit estimation results for the effect of two labor market variables on coresidence: an indicator for whether the youth was working in a given month and log monthly earnings. For the models with log earnings as the independent variable, the sample is restricted to working youths. All models include a large set of fixed and time-varying control variables (see table note for details). Columns 1-4 show that working youths are significantly more likely to be living away from their parents than are nonworking youths and that among working youths, those with higher earnings are also more likely to live away from home. When lagged coresidence is included as a control, the results are stronger, suggesting that dynamics may play an important role in the analysis. To control for unobserved heterogeneity across youths in their propensity to live away from home that may be correlated with labor market outcomes, columns 5 and 6 report results from a fixed-effects (conditional) logit. The results with fixed effects are not statistically significant but still indicate a positive relationship.

TABLE 2
Proportional Hazard Models for Coresidence Dynamics

	Moving Back with Parents			Moving Out Again	
	(1)	(2)	(3)	(4)	(5)
Working	.689***			1.274	
	(.088)			(.353)	
Recently stopped work[a]		1.628***			
		(.249)			
Log earnings			.888***		1.842***
			(.034)		(.380)
Includes random effects	Yes	Yes	Yes	Yes	Yes
Includes duration controls				Yes	Yes
Observations	10,330	10,330	8,041	2,241	1,569
Individuals	779	779	704	267	225

Note.—Conditional log-log models for discrete time proportional hazard. See App. B for full details. All models include controls for race, high school graduation, marriage, biological children, parental education, parental income, biological parents married, cubic polynomial in youth age, quadratic polynomial in average parent age, and region dummies, as well as normally distributed random effects (frailty). Standard errors are in parentheses, clustered by individual. Reported coefficients are multiplicative effects on the proportional baseline discrete time hazard.

*** Significant at 1 percent level.

[a] Equal to 1 if a youth transitioned from employment to nonemployment in the current or preceding 3 months.

Duration Analysis

The above logit analysis is useful for understanding coresidence at a point in time. However for high-frequency data with infrequent transitions across states, an explicit duration analysis for the hazard rate of moving in and out of home is more appropriate. I employ a discrete time proportional hazard model with random effects (frailty) to estimate how the probability of transitioning across coresidence states is affected by youths' current and recent labor market outcomes. The random effects control for unobserved individual-specific heterogeneity in the baseline hazard of moving in or out of home. This duration-based approach improves on the preceding analysis (and existing empirical studies in Rosenzweig and Wolpin [1993, 1994]; Ermisch and Di Salvo [1997]) since it more fully exploits the high-frequency dynamics of parental coresidence by using information on labor market conditions at the time that transitions take place. Full details of the econometric model are in Appendix B.

The results are shown in table 2. Columns 1–3 show how the probability of moving back home is affected by three measures of a youth's labor market situation. A youth who is currently working is around 31 percent less likely to move back home than a similar youth who is out of work. A youth who has recently stopped working (defined as transiting from employment to nonemployment in the previous 3 months) is around 63 percent more likely to move back home than is a youth who has not undergone such a transition. For working youths, reductions in earnings are also associated with an increased likelihood of moving back home: a 10 percent decline in earnings increases the hazard of moving back by around 1 percent. Columns 4 and 5 show how the probability of moving out again, after having moved back home, responds to labor market outcomes. The models for moving out of home again include controls for duration dependence in the hazard rate and spell-specific random effects. The standard errors are larger because of the relatively small number of spells back home, but the point estimates are large. Being employed is associated with a 27 percent increase in the likelihood of moving out again, and a 10 percent increase in earnings increases the likelihood by around 8 percent.

Longer-Run Effects of Coresidence

The evidence in the previous two tables suggests that coresidence is potentially important for smoothing short-term labor market fluctuations. I now provide evidence that coresidence, and the ability to transit coresidence states, is also potentially important for long-term labor market outcomes. In particular, I show that job separations have lasting effects on youths' earnings but that these effects are strongly mitigated through coresidence and the ability to move back home after a job separation.

To illustrate this, I make use of more recent waves of the NLSY97 to look at earnings at age 26. I compare these earnings to coresidence and labor market dynamics 6 years earlier, at age 20. For the subsample of youths who were working at age 20, I regress log earnings at age 26 on log earnings at age 20 and an indicator variable for whether the youth underwent a transition from employment to nonemployment while aged 20.

The results are displayed in table 3. Panel A reports results without any additional control variables, while panel B reports results when a large set of individual-level controls are included in the regression. Focusing on the results that include control variables, column 1 shows that a job separation at age 20 can have a significant long-term effect on earnings. Youths who experienced such a separation have earnings at age 26 that are on average 25 percent lower than youths who did not. In columns 2 and 3, I split the sample into youths who were living away from their parents at the time of the separation and youths who were living with their parents, respectively. The results show that all of the negative long-term effect of the job separation is concentrated among those youths who were living away from home. For such youths,

TABLE 3
Longer-Run Impact of Job Loss and Coresidence
(Dependent Variable: Log Earnings at Age 26)

	Ordinary Least Squares				
	(1)[a]	(2)[b]	(3)[c]	(4)[d]	(5)[e]
	A. Without Additional Controls				
Stopped working at age 20	−.286**	−.454***	−.189	−.346	−.428*
	(.126)	(.171)	(.163)	(.253)	(.247)
Log earnings at age 20	.228***	.220**	.226***	.173	.265***
	(.066)	(.087)	(.084)	(.178)	(.090)
	B. Including Additional Controls				
Stopped working at age 20	−.251*	−.635***	−.129	.483	−.380**
	(.128)	(.179)	(.162)	(1.062)	(.187)
Log earnings at age 20	.208***	.297***	.209***	−.254	.275**
	(.064)	(.104)	(.078)	(.458)	(.122)
Individuals	349	115	234	28	87

Note.—Additional controls in panel B include race, high-school graduation, marriage, biological children, parental education, parental income, biological parents married, cubic polynomial in youth age, quadratic polynomial in average parent age, and region dummies.
* Significant at 10 percent level.
** Significant at 5 percent level.
*** Significant at 1 percent level.
[a] All working youths at age 20.
[b] Youths living away from their parents at the time of stopping work.
[c] Youths living with their parents at the time of stopping work.
[d] Youths living away and who moved back in after stopping work.
[e] Youths living away and who did not move back in after stopping work.

the long-term effect is very large, around 64 percent, whereas for youths who were living with their parents at the time of the separation, there was no significant effect on long-term earnings (and the point estimates are negative). In columns 4 and 5, I further split the sample of youths that were living away from their parents into those who moved back home in the 3 months after stopping work and those who did not, respectively. On average, those youths who did not move back home suffered significant long-term earnings losses by age 26, while the effect for those youths who did move back home is much smaller and not significantly different from zero (of course, when studying such an extreme subsample, the number of observations necessarily becomes very small, motivating the need to substantiate this evidence with a structural model).

B. *Repeated Cross Sections from CPS*

The rich individual-level data and panel structure of the NLSY97 make it ideal for estimating the effects of labor market variables on coresidence transitions at the individual level. However, there are two important drawbacks of those data. First, observed labor market changes reflect both exogenous and endogenous changes, including some labor market transitions that may be in anticipation of a move in or out of home. Second, the data structure necessitated a focus on males with a high school education or less. In this section, I exploit repeated cross sections from the CPS to show that the link between parental coresidence and employment holds for a much larger part of the population. I start by documenting time series variation in parental coresidence at business cycle frequencies to establish that the coresidence margin is quantitatively important in the aggregate. I then use a difference-in-difference identification strategy that exploits differential labor market conditions across US states to suggest a causal relationship between the labor market and coresidence rates. This identification strategy is a useful complement to the one based on individual-level variation and fixed effects used in Section II.A.[8] Details of the data can be found in Appendix A, Section B.

Time Series for the United States

The plots in figure 1 show the cyclical components of the time series for the parental coresidence rate and market hours, for the populations aged 16–24 and 25–34. For both groups, these plots reveal a strong time series correlation between coresidence and hours worked. This correlation is most pronounced during the large economic downturns of 1982 and 2008. It suggests that when labor market conditions worsen, more people live with their parents. Hence, even though the structural model and NLSY97 data focus on very young low-skilled youths only, these plots suggest that the coresidence channel is likely to be important for other education groups and for older workers: the CPS data contain individuals of all education levels, and the figure displays strong comovement, even among people aged 25–34.[9]

Cross-State Variation

Further evidence for the linkage between coresidence transitions and the labor market can be obtained by exploiting cross-state variation in

[8] I thank an anonymous referee for this suggestion.

[9] Motivated by this time series evidence, Dyrda, Kaplan, and Rios-Rull (2012) embed a coresidence decision into a real business cycle model that admits aggregation, and they show that endogenous household formation is quantitatively important for aggregate labor market fluctuations. The relevance of the coresidence margin depends crucially on the parameters that govern decisions about when to move back home and further underscores the value of estimating structural parameters using individual-level panel data.

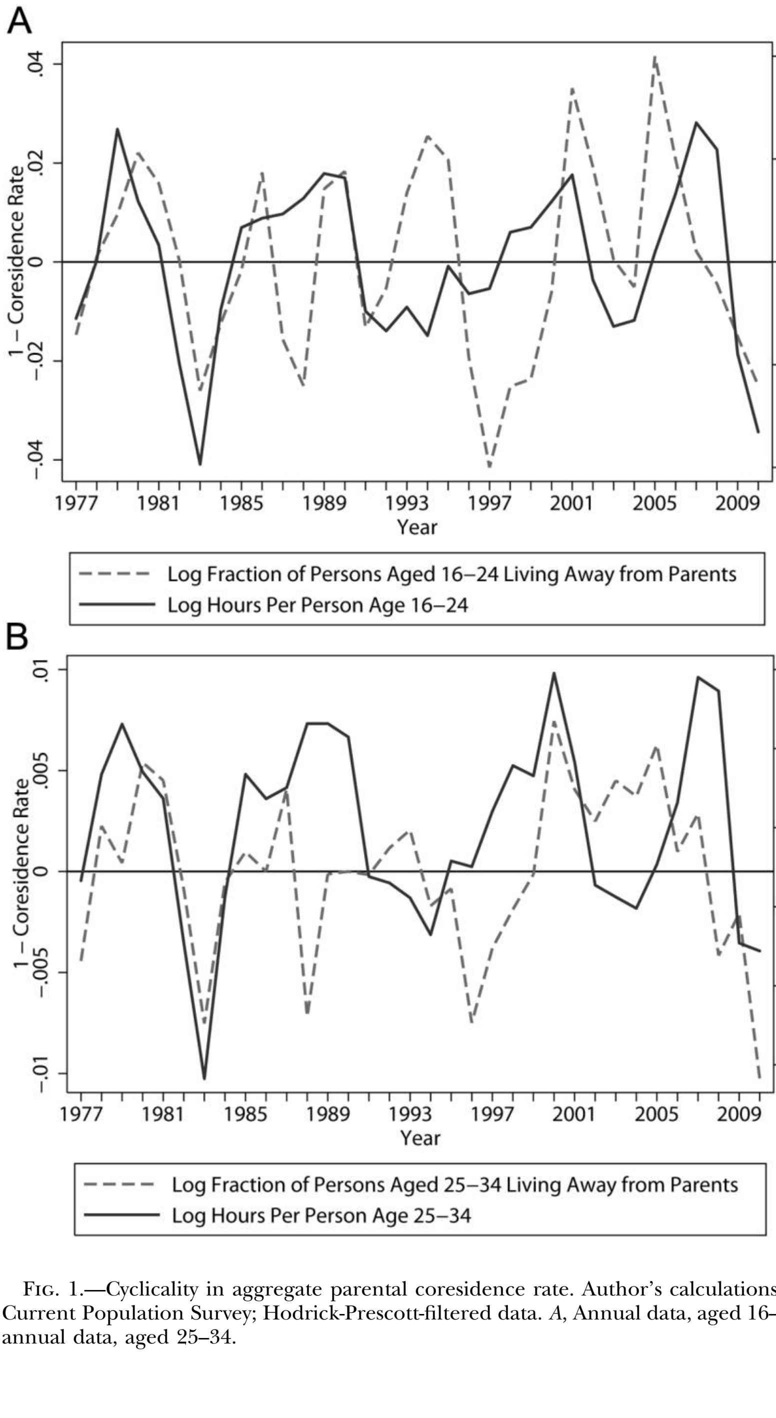

Fig. 1.—Cyclicality in aggregate parental coresidence rate. Author's calculations from Current Population Survey; Hodrick-Prescott-filtered data. *A*, Annual data, aged 16–24; *B*, annual data, aged 25–34.

TABLE 4
EFFECT OF LABOR MARKET VARIABLES ON CORESIDENCE RATES ACROSS US STATES

	MONTHLY DATA			QUARTERLY DATA		
	(1)	(2)	(3)	(4)	(5)	(6)
Employment rate:						
16–34	−.279***	−.218***	−.151*	−.257***	−.211***	−.137*
	(.061)	(.056)	(.058)	(.061)	(.055)	(.063)
16–24	−.304***	−.105*	−.131*	−.294***	−.083	−.107
	(.056)	(.050)	(.054)	(.056)	(.050)	(.061)
Hours worked:						
16–34	−.183***	−.130***	−.098***	−.186***	−.132***	−.098**
	(.031)	(.031)	(.026)	(.031)	(.031)	(.029)
16–24	−.142***	−.049*	−.081***	−.148***	−.044	−.078***
	(.025)	(.024)	(.019)	(.025)	(.024)	(.021)
State fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Varying state controls	No	Yes	Yes	No	Yes	Yes
State-specific trends	No	No	Yes	No	No	Yes

NOTE.—Cumulative effect over 2 years of a 1-percentage-point increase in the employment rate, or a 1 percent increase in average hours worked for the relevant age group. Data cover 1979–2010. Time-varying controls include average age and Housing Price Index. State time averages are weighted by population. Standard errors are in parentheses, clustered by individual.
* Significant at 10 percent level.
** Significant at 5 percent level.
*** Significant at 1 percent level.

aggregate labor market conditions and coresidence rates from the CPS. I use monthly and quarterly data from 1979 to 2010. I consider two age groups: a narrow range from 16 to 24 to be consistent with the NLSY97 data and the scope of the structural model and a broader range from 16 to 34 to show that the effects extend to individuals in their late twenties and thirties. I regress the coresidence rate in each state and time period on each of two measures of the current labor market conditions in that state: the employment rate of the relevant age group and the log of average hours worked by that age group. Each regression includes a full set of state and time fixed effects, so that identification is effectively through differences-in-differences, as well as 2 years worth of lagged labor market variables to allow for dynamics in the adjustment of living arrangements.

The baseline results are shown in column 1 of table 4 for monthly data and column 4 of table 4 for quarterly data. Each number reflects the cumulative effect over 2 years of a 1-percentage-point increase in the employment rate, or a 1 percent increase in average hours worked, on the fraction of people in each age group that live with a parent. The results show statistically significant and economically substantial effects of labor market conditions on coresidence rates. For each percentage-point increase in the employment rate, the coresidence rate increases by between 0.25 and 0.3 of a percentage point, while a 1 percent increase in average hours worked leads to a cumulative increase in coresidence of 0.15 to 0.18 of a percentage point.

The remaining columns in table 4 show how these results change when additional controls are included. Columns 2 and 5 add two time-varying state-specific controls that are likely to be correlated with both the labor market and coresidence: the average age in each age group and the Housing Price Index from the Federal Housing Finance Agency. In addition, columns 3 and 6 control for state-specific trends in coresidence rates. Including these additional controls reduces the quantitative magnitudes of the estimated coefficients, but all remain economically large and statistically significant.

The Need for a Structural Model

At this point it is worth taking stock of what we can learn from this empirical evidence. The data do indeed provide evidence that (i) there is a causal link between labor market outcomes and the propensity for youths to move in and out of their parents' homes and (ii) the option to move back home in response to shocks has a significant impact on subsequent labor market outcomes. However, without a structural model, the inferences that can be drawn are still limited. There are a number of reasons.

First, given the high-frequency panel data required for the analyses above, the available sample sizes are small, particularly after conditioning on youths who have experienced particular labor market and coresidence histories as in the analysis of Section II.A. Second, in the data, one can only identify youths who actually move in with their parents, as opposed to youths who have the option to move in with their parents. Moreover, it is difficult to think of a natural experiment or potential instrumental variable that could stand in for this, which limits the feasibility of a purely empirical approach. In a model, however, it is possible to directly control this option. Thus, a structural model makes it possible to quantify the importance of the option to move back home both on short-term consumption smoothing (i.e., insurance) and longer-term outcomes that operate through labor market decisions. In Section V, I use realized preference shocks, which are observed in the model and are orthogonal to labor market shocks, as a proxy for this option. Third, without a model, one cannot consider counterfactual environments, such as removing the option to live at home, or policy proposals that affect incentives to live with parents. In the remainder of the paper, I develop a model that is rich enough to achieve all these goals yet places enough structure on the data to enable estimation of the key parameters governing coresidence, labor supply, and transfer decisions.

III. Model

A. *Environment*

Demographics

Time is discrete and measured in months. I focus on the finite horizon $t = 0, 1, \ldots T$. The basic unit in the model is a family, which consists of a parent-youth pair. Families are indexed by i, and each comprises a parent (p) and a youth (y). In any month, t, the family can be in one of two residential states, labeled $r_{it} \in \{0, 1\}$. When $r_{it} = 0$, the youth lives in the parental home, and when $r_{it} = 1$, the youth lives in separate housing away from his parents.

Youth Preferences

Youths have time-separable, expected-utility preferences, defined over consumption, labor supply, and their residence state. Let U^y_{it} denote the period utility for a youth from family i:

$$U^y_{it} = \frac{[c^{y(1-\phi)}_{it} G^{\phi}_{it}]^{1-\gamma}}{1-\gamma} - h_{it}v + r_{it}z_{it}, \tag{1}$$

$$G_{it} \equiv g^y_{it} + (1 - r_{it})g^p_{it}. \tag{2}$$

Period utility is additively separable between utility from consumption, labor supply, and direct utility from independence. There are two types of consumption goods: c^y_{it} is the youth's consumption of a private good, enjoyed exclusively by the youth; G_{it} is total consumption of a locally public good inside the home. It consists of the youth's own purchases of the good, g^y_{it}, as well as public consumption purchased by the parent that is available inside the parental home, g^p_{it}. Of course, the youth only has access to this second quantity of the public good if he lives at home (i.e., $r_{it} = 0$).

In reality, most types of consumption fall somewhere in between these two extremes: a nonrival and nonexcludable good (g) and a completely private good (c). The extent to which there exists economies of scale in the parental home, or in other words, the extent to which most consumption is like g rather than c, is a key determinant of the value of coresidence and the impact of coresidence on behavior. The Cobb-Douglas specification allows for a parameter ϕ (the weight on public goods in the utility function) that indexes this extent. When $\phi = 1$, only public goods are consumed, and there are full economies of scale;

when $\phi = 0$, all consumption is private, and there are no economies of scale. In Section IV.A, I use data on consumption equivalence scales to calibrate a value for ϕ. I assume a constant relative risk aversion utility function over the composite consumption good, with risk aversion γ.

The direct utility from living away from home is denoted by z_{it}. It is stochastic and differs across youths. In any month, a youth can be either working, $h_{it} = 1$, or not working, $h_{it} = 0$. The disutility of working is constant and fixed at v. Lifetime utility for a youth is given by

$$V_0^y = E_0\left\{\sum_{t=0}^{T} \beta^t U_{it}^y + \beta^{T+1} V_{T+1}^y\right\}, \tag{3}$$

where V_{T+1}^y is a terminal value function described below.

Preference shocks, z_{it}, are assumed to follow a discrete-state Markov process whose variance, σ_z^2, and autocorrelation, ρ_z, are constant with age. The preference shocks play an important role in the model and should be interpreted as a reduced-form way of capturing the effects of non–labor market heterogeneity in the relative preference for living away from home. Such shocks may include the formation and dissolution of cohabiting relationships, peer effects, and changes in the demographic structure of the parental home. In reality, these effects are likely to exhibit an increasing trend with age, making living away from home an increasingly attractive option for young adults as they get older. For example, independence from one's parents is itself something that becomes more attractive as youths move toward adulthood. To capture this feature of coresidence, the mean relative preference for living away from home, $E[z_t] = \alpha_z + \beta_z t$, is allowed to increase exogenously according to a linear trend.

It will become apparent that the model features a number of endogenous mechanisms for generating the observed increase in the fraction of youths living away from home between ages 17 and 22, all related to the labor market experience of youths: (i) an increasing probability of working, (ii) an increasing earnings profile conditional on working, and (iii) asset accumulation. If $E[z_t]$ were assumed to remain constant with age, the model would risk assigning an overly important role to these factors in determining coresidence patterns. By allowing for flexibility in the mean growth rate, variance, and autocorrelation of z, the model is such that in principle, either labor market or non–labor market factors could be the primary driver of coresidence outcomes. The features of the data that help identify these parameters, and hence distinguish between these two hypotheses, are discussed in Section IV.

Parent Preferences

Parents have time-separable expected-utility preferences and are altruistic toward their children. They have direct preferences over their own private consumption, c_{it}^{p}, and public consumption, g_{it}^{p}:

$$U_{it}^{p} = \frac{[c_{it}^{p(1-\phi)} g_{it}^{p\phi}]^{1-\gamma}}{1-\gamma}. \tag{4}$$

Their total utility, V_0^p, consists of their direct utility, $\tilde{V}_0^p$, plus the utility of their child, V_0^y, weighted with an altruism factor, $\eta \geq 0$:

$$V_0^p = \tilde{V}_0^p + \eta V_0^y,$$

$$\tilde{V}_0^p = E_0\left\{\sum_{t=0}^{T} \beta^t U_{it}^p + \beta^{T+1} V_{T+1}^p\right\}.$$

These preferences imply that parents do not get utility from public consumption purchased by the youth (g_{jt}^{y}) when the youth lives at home. This asymmetry is unlikely to have any meaningful effect on allocations because in most cases the youth does not make additional purchases of public consumption when living at home.

This form of one-sided altruism has a long history in the modeling of parent-child interactions and has a number of implications for behavior (see, e.g., Altonji, Hayashi, and Kotlikoff [1997] and references therein). First, altruism is the mechanism that is used to generate financial transfers from parents to youths. Second, note that when $z_{it} >$ 0, both parents and youths have a preference for the youth to live away. However, since in general $\eta < 1$, parents have a weaker direct preference for youths to live away than do youths themselves. This conflict may manifest itself as multiple equilibria of a simultaneous-move version of the game described below and motivates the need to specify an appropriate timing protocol for the game. Altruism also implies that parents have a stronger preference for youths to work at a given wage, generating a second form of conflict. Both forms of conflict can generate inefficiencies, which are discussed in the context of a static version of the game that admits a closed-form solution in Appendix C.

Budget Constraints

In each period, a youth can be in one of two labor market states: employed ($h_{it} = 1$) or unemployed ($h_{it} = 0$). An employed youth earns an idiosyncratic monthly wage w_{it}, which is the outcome of a stochastic process, outlined in the section below on the labor market. An unemployed youth receives an exogenous amount b. This should be interpreted as the benefit from a simple public unemployment insurance program. Labor income taxes are levied according to the function τ.[10]

Youths can use their income to purchase private and public consumption goods, c^y_{it} and g^y_{it}, and to invest in a risk-free asset, $a_{i,t+1}$, which earns interest at a gross rate R. Borrowing is not allowed, so $a_{i,t+1} \geq 0$. In addition, a youth may receive a transfer $T_{it} \geq 0$ from his parents. There is a per-period fixed monthly cost of housing, χ, payable by youths living away from home and a fixed cost κ of moving out of home. The per-period cost is intended to capture both direct housing costs such as rent and mortgage payments as well as indirect costs such as gas and electricity bills. The fixed cost of moving out is intended to capture direct moving costs as well as indirect costs that may include purchases of new furniture and other durable consumption goods. There is no cost to move back home. The youth budget constraint is hence given by

$$\begin{aligned} & c^y_{it} + g^y_{it} + a_{i,t+1} + r_{it}[\chi + (1 - r_{i,t-1})\kappa] \\ & \leq w_{it}h_{it} - \tau(w_{it}h_{it}) + b(1 - h_{it}) + Ra_{it} + T_{it}. \qquad (5) \end{aligned}$$

Parents have an exogenous constant income stream, I^p_i, which differs across families. Parental income can be used to purchase private consumption goods, c^p_{it}, and public consumption goods, g^p_{it}. In addition, parents can make nonnegative financial transfers, T_{it}, to youths. The parental budget constraint is hence given by

$$c^p_{it} + g^p_{it} + T_{it} \leq I^p_i - \tau(I^p_i) - \chi. \qquad (6)$$

Resource Sharing across Generations

To understand the mechanisms in the model for sharing resources across generations, it is worth taking stock of the various technologies for transferring utility from parents to youths. The model features two forms of parental support: coresidence and financial transfers. For a youth living away from his parents, financial transfers are the only means that parents have to share resources: providing an additional unit of assets to the youth requires the parent to forgo one unit of current period consumption.[11]

Coresidence can also be interpreted as a technology for intergener-

[10] The monthly wages from the NLSY97 that are used in the estimation of the model are gross of labor income taxes, which necessitates the inclusion of a tax function in the model. The assumed tax function is based on the US tax system in 2007 and is described in App. F.

[11] Because a nonresident youth's savings or labor supply decision may adjust in response to a transfer from the parent, the cost to the parent of providing a unit of consumption, rather than assets, to a nonresident youth may be different from one. The actual cost is given by $\partial c^y/\partial a$, where a is the resources of the youth.

ational transfers. Using it incurs a fixed monthly cost through forgone utility from independence (z for the youth, η for the parent). The return from paying this utility cost comprises two parts. First, there is a fixed monthly benefit from the savings in housing costs (χ). Second, coresidence reduces the cost of transferring additional units of resources from parents to youths, through the presence of the public good. The extent of this "cost-saving" benefit of coresidence is determined by the weight on public goods in the utility function, ϕ, which in this setting plays the role of economies of scale. Conditional on living at home, the marginal cost for the parent of providing an additional unit of resources to the youth is decreasing in ϕ and approaches zero as ϕ approaches one (full economies of scale).

Note also that the only component of coresidence that does not enter through consumption is the forgone utility from independence when living at home (z). With decreasing marginal utility, this component dominates as consumption increases, implying that z becomes relatively more important (and hence coresidence less attractive) as the youth's assets increase. It also implies that as assets increase, preference shocks become relatively more important, and labor market shocks become relatively less important, for generating coresidence dynamics.

Labor Market Search

The labor market is a frictional environment consisting of two types of jobs that I denote as risky (R) and safe (S). Frictions take the form of the random arrival of job offers plus job destruction shocks and wage shocks while working. Job offers can be refused, and youths always have the option of quitting employment. Employment status is denoted by $h_{it} \in \{0, 1\}$, and job types are denoted by $j_{it} \in \{S, R\}$. In addition to having initial offers drawn from separate distributions, the two types of jobs differ in terms of their riskiness and potential for future earnings growth. I defer my discussion of the reasons for modeling jobs in this way until after they have been described.

Youth wages consist of two components that are additive in logs:

$$\log w_{it} = \vartheta_t + \omega_{it},$$

where ϑ_t is a deterministic age effect that reflects general experience effects, assumed to grow linearly at rate μ_ϑ, and ω_{it} is an individual-specific and job-specific component that is the outcome of the following random search process.

All labor market shocks are assumed to be realized at the beginning of each period. At the beginning of month t, a youth who was not working in month $t-1$ receives an offer to work at a type j job with probability λ_0^j. Offers are drawn from a lognormal distribution:

$$\omega_{it} \sim N(\mu_0^j, \sigma_0).$$

Note that the mean of the offer distribution differs by job type. A youth who receives an offer may accept it and work in period t or reject it and hope to receive another offer in month $t+1$. Offers from type $j = S$ jobs remain constant at their initial value for the duration of the job spell. This is the sense in which these are safe jobs. With probability δ, the job (regardless of type) is exogenously destroyed. A youth who loses his job in this way must spend period t not working, and so has $h_{it} = 0$.

At the beginning of month t, a youth who worked in month $t-1$ at a wage $w_{i,t-1}$ in a type $j_{i,t-1} = R$ job receives a new wage draw. New wages are assumed to follow a random walk in logs with drift:

$$\log \omega_{it} = \mu_d + \log \omega_{i,t-1} + \varepsilon_{it},$$

$$\varepsilon_{it} \sim N(0, \sigma_1).$$

The youth has the option of rejecting the new wage offer in favor of nonemployment but does not have the option of staying at his current job at his existing wage. The drift μ_d will be estimated to be positive, implying that on average, on-the-job wage shocks represent good news, and R-type jobs yield upward-sloping expected earnings profiles. These wage shocks with drift are intended to be a reduced-form way of capturing fluctuations in the quality of worker-firm matches, new job offers that arise from on-the-job search, and returns to tenure. However, wage changes may be negative, yielding new wages that are below existing wages. This is the sense in which these are risky jobs.

The labor market is modeled in this way because it captures the notion that short-term insurance possibilities, including coresidence, may have important longer-term consequences for labor market outcomes through the job selection process. Part of the reason is that, as with all search models, there is an option value to waiting for better offers, so that youths who can afford to wait longer for jobs will earn higher wages. To the extent that jobs persist, this generates better long-term outcomes.[12]

This is true for both types of jobs in this economy. However, since the estimates will imply that $\lambda_0^S > \lambda_0^R$, so that S-type jobs arrive more frequently than R-type jobs, but R-type jobs deliver expected earnings growth, there is an additional benefit of being able to wait for offers: the jobs that offer the best long-term prospects require youths to wait longer on average. Hence, R-type jobs are more attractive to youths with good access to resources while unemployed. Indirect evidence for this type of mechanism was provided in Section II.A. Finally, there is the

[12] Acemoglu and Shimer (1999) exploit this logic as a motive for providing unemployment benefits to liquidity-constrained unemployed workers.

effect of risk and its relationship to future insurance possibilities. Since R-type jobs are both those with higher earnings growth and those that feature earnings uncertainty, youths will be more inclined to wait for R-type jobs if they have better access to future insurance (e.g., are less averse to moving back home should the need arise). How youths trade off these motives depends on their insurance possibilities—assets, parental income, and disutility of coresidence—hence, the equilibrium of the game between youths and parents features reservation wages that are a function of all these state variables.

Finally, I note that using data on wages and employment alone, a negative wage shock combined with a quit cannot be empirically distinguished from job destruction. Together with the fact that rejected offers to the unemployed are not observed, this implies that the labor market parameters cannot be estimated outside the model in a first stage and fed into the structural model. Instead they must be estimated along with the other structural parameters inside the model. The parametric assumptions on the distribution of shocks help to achieve identification. Identification of the labor market and other parameters is discussed further in Section IV.

Government Insurance

In the NSLY97 sample described in Section II, 17 percent of male youths received a government benefit at some point during their time in the sample. With the exception of unemployment benefits that are received by 6 percent of the sample, these are all means-tested benefits.[13] Following Hubbard et al. (1995), I model means-tested benefits as a consumption floor, $\underline{c}$. Hubbard et al. (1995) show that allowing for the effects of means-tested benefits is important in understanding savings behavior of poor households. Unemployment benefits, b, are modeled as a constant benefit that is paid automatically to youths in any period that they are not working. The tax function, described in Appendix F, is progressive and as such is an additional form of government insurance.

Initial Conditions and Terminal Values

A complete description of the model also requires specification of initial conditions for assets, residence, and labor market variables. Youths are assumed to have $a_0 = 0$ at age 17. This a reasonable assumption, given

[13] Means-tested benefits include food stamps, Aid to Families with Dependent Children (AFDC), and Women, Infants, Children (WIC). WIC is the most common (received by 7.5 percent of the sample). This is surprising because it is intended only for females; 90 percent of these recipients are in a cohabiting relationship, suggesting that respondents report benefits received by all members of the household.

the large fraction of youths with exactly zero assets in the NLSY97 data.[14] All youths are assumed to be living at home at $t = -1$. This assumption only places a minimal exogenous structure on r_0, which is due to the fixed cost of moving out. Since r_0 is itself a choice variable, youths are free to move out in the first model period. An exogenous fraction of youths are assumed to have been working at $t = -1$, all in S-type jobs. Their wages are given by the observed distribution of monthly earnings at age 17. Again, these youths are free to quit to unemployment immediately at $t = 0$.

Because of the monthly frequency of the model, it is not computationally feasible to solve and estimate the model using a horizon T that corresponds to the end of the life cycle. Moreover, because our interest is in producing a good model of high-frequency behavior around the time of entry to the labor market, it is not clear that this would be a preferred approach even if it were computationally feasible, given the inherent danger of misspecification in any model. Instead, I choose to specify terminal value functions and solve backward from these. In order to minimize the impact of assumptions about functional form, I solve the model for an additional 2.5 years (30 periods) past the point at which I have data (age 22). I then compare the model and data over the 6 years covering ages 17–22.

The assumption at $t = T$ is that the interaction between parents and youths ceases, and no more financial transfers can be made. At this point, all youths that are still living at home are forced to move out, labor supply becomes inelastic, and there is no further uncertainty about future wages. These assumptions are sufficient to obtain closed-form solutions for the value functions, which are then used as the terminal values.[15]

Feasible Allocations

Consider stochastic processes for labor market variables (ω_t, j_t) and preference shocks (z_t, where i subscripts have been dropped to ease nota-

[14] Excluding cars, only 21 percent of 17-year-old males have positive wealth, and 75 percent have exactly zero wealth. If cars are included, then 44 percent have positive wealth, and 55 percent have exactly zero wealth.

[15] All results are unchanged when the model is solved with an additional 10 years (120 months). Also, note that because of the low estimated discount factor for this group of low-skilled males, the effect of the form of the terminal value assumption on parameter estimates is minimal. An alternative approach would be to specify the terminal value functions as unknown parametric functions of the state variables and to estimate these functions along with the other structural parameters. For this approach to be feasible, it is necessary to have high-quality data on the state variables in the final period (including assets). However, after age 20, asset information in recent waves of the NLSY97 is only collected every 5 years and, hence, is only available for a subset of the sample at age 22. Moreover, even with this asset information, it is unlikely that identification of the terminal value functions could be achieved without additional assumptions about functional forms.

tion); an exogenously given interest rate, R; and cost of housing, χ. An allocation in this environment, s, is a sequence of functions that map histories of labor market outcomes and preference shocks $\{\omega^t, j^t, z^t\}$, initial conditions $\{a_0, w_{-1}, h_{-1}, j_{-1}, r_{-1}\}$, and heterogeneity in parental income $\{I^p\}$ into values for $\{r_t, h_t, c_t^p, c_t^y, g_t^p, g_t^y, T_t, a_{t+1}\}$. An allocation is feasible if it satisfies the parental budget constraint (6), the youth budget constraint (5), and the nonnegativity constraints for transfers, assets, and consumption. Denote the set of feasible allocations by $\mathcal{S}$.

Markov-Perfect Equilibrium

To determine allocations in this environment, I adopt a decentralized approach without commitment, whereby an equilibrium concept and a timing protocol are specified, and parents and youths make strategic decisions to maximize their expected discounted lifetime welfare. The environment just described has a natural interpretation as a stochastic repeated game in which action sets in the stage game are conditioned on a payoff-relevant state vector consisting of the current asset position, a_t; the residence state in the previous period, r_{t-1}; and the realized values of the shocks (w_t, z_t, j_t). In each repetition of the stage game, the youth chooses whether to reside at home or away, whether to work, and how much to save and purchase of each type of consumption good. The parent chooses monetary transfers to be paid to the youth and how to split his or her income between public and private consumption. The distribution of $(w_{t+1}, z_{t+1}, j_{t+1})$ is determined by (h_t, w_t, z_t, j_t). The equilibrium concept that I propose is an MPE in which all actions are conditioned on payoff-relevant variables only; $x_t = (a_t, r_{t-1}, w_t, z_t, j_t, I^p)$.[16]

Timing of Stage Game

In order to guarantee uniqueness of the MPE, I impose a particular extensive form of the stage game that specifies the order in which parents and youths make their decisions. Attention is then restricted to the unique subgame perfect equilibrium of this sequential stage game. The

[16] In App. E, I describe an alternate approach for determining allocations in which there is full commitment. In that environment, attention is restricted to the subset of $\mathcal{S}$ that is Pareto efficient between youths and parents, taking prices and stochastic processes as given. There I give arguments to support the decentralized approach without commitment as the preferred modeling choice, despite it being computationally more challenging. Due to the finite horizon and the fact that the adopted timing protocol guarantees a unique subgame perfect equilibrium in the stage game, the unique MPE is also the unique subgame perfect equilibrium of the dynamic game. In an earlier working paper (Kaplan 2010), I show that at the estimated parameters, the values from the game lie very close to the Pareto frontier. Hence, considering alternative timing protocols might lead to different equilibria but not to much more efficient ones.

assumed timing is as follows. First, the current state $x_t = (a_t, r_{t-1}, w_t, z_t, j_t)$ is observed. Then the youth chooses whether to live at home or away (r_t). Next, the parent chooses monetary transfers, T_t, and public consumption (g_t^y). Finally, the youth makes his current period labor supply, h_t, and consumption-savings decision (c_t^y, g_t^y, a_{t+1}).

The reason for specifying a timing protocol for the stage game is that the simultaneous-move version may contain multiple Nash equilibria. This is most easily demonstrated in the one-shot static version of the game with exogenous labor supply in Appendix C. The intuition for the appearance of multiple equilibria is that due to imperfect altruism ($\eta < 1$), parents have a weaker direct preference for the youth to live away from home than does the youth. This generates values for preferences, z, such that the youth prefers an equilibrium in which he lives away and receives the resulting optimal transfer, while the parent prefers an equilibrium in which the youth is induced to stay at home by the (noncredible) threat of low transfers if he were to move out. However, with the assumed timing protocol, these latter equilibria are not subgame perfect and are ruled out.

This particular timing protocol is motivated in part by casual observation of the way that these interactions take place in reality. It seems reasonable that parents cannot force youths to adhere to a particular consumption/savings policy or labor supply decision rule. Rather, they can only influence these choices through their choice of financial transfers. Similarly, a youth cannot be forced to stay in the parental home if he wants to move out, but he must accept whatever resulting transfer the parent decides to make. What he does with that transfer is up to him.

B. *Determinants of Coresidence and Transfers*

In Appendix D, I show how the MPE can be described by a set of Bellman equations, and I provide a formal description of the decision problems of parents and youths along the equilibrium path. Since there is no analytic solution to this game, the MPE must be computed numerically by backward induction, and a formal characterization is not possible. Nonetheless, a number of features of the determinants of transfers and coresidence can be described qualitatively.

There are four model elements that contribute to an increasing fraction of youths living away from home with age: (i) an increasing average earnings profile, (ii) an increasing probability of employment through search, (iii) asset accumulation, and (iv) an increasing preference for living away from home. Youths are more likely to live away from home when earnings, assets, or the value of independence is higher. However, the probability of living away from home is ambiguous with respect to parental income. On the one hand, higher parental income generates higher parental transfers and hence a lower earnings/assets threshold for the youth to live away. On the other hand, higher parental income means higher consumption in the parental home, making living at home a more attractive option for the youth.

There are two classes of reasons why a youth may move back home: (i) economic factors that include job loss, wage drops, lower than expected earnings growth, and asset decumulation and (ii) preference shocks (z). However, the dynamics of coresidence outcomes for these two types of shocks are likely to be different. Because of the ability to run down assets, a labor market shock may lead to a move back in a subsequent period, rather than in the same period in which the shock occurred. However, preference shocks, if they lead to a move back home, are likely to do so in the current period, provided the shocks are sufficiently persistent (which is the case for the estimated parameter values).

Transfers are determined by equating their marginal value and marginal cost to the parent along the equilibrium path. When the youth lives away from home, the marginal cost of a transfer is the value of current period consumption for parents. The marginal benefit comprises two parts. The first part is due to altruism: the marginal value of assets for youths, scaled by the altruism factor. The second part is the marginal benefit to parents of the youth holding more assets, which accrues from lower expected transfers in the future. When the youth lives at home, the parent's problem is complicated by the fact that the youth benefits from funds spent on the public consumption good by parents. However, if the youth's optimal choice between public and private consumption (c_t^y, g_t^y) is at an interior point, then it is always optimal for the parent to set $T_t = 0$. To see this, note that the parent could always obtain higher utility by spending marginal funds on public consumption, g_t^p, rather than transfers, T_t. That way, the parent gets positive utility from additional current consumption, while the youth can cut back on his own expenditures on g_t^y and so is unaffected. In other words, parental spending on public consumption crowds out the youth's spending on current consumption one for one. However, if the youth's optimal choice for public consumption is at the corner ($g_t^y = 0$), then it may be optimal for the parent to make positive financial transfers even when the youth lives at home.

C. *Discussion of Modeling Assumptions*

Before describing the estimation procedure for the structural parameters, there are three features of the model that warrant further discussion.

Two Residence States

The model assumes that there is only one home state and one away state. Implicitly this means that the model ignores the possibility for a youth who is not living with his parents to cohabit with other people, such as a partner, siblings, or roommates. The NLSY97 contains information on the household composition of youths that can be used to assess the restrictiveness of this assumption. An important caveat is that this information is available only on the annual interview date and only for the household that the youth considers to be his permanent residence at that time. In many cases this household does not correspond to the youth's actual living situation. To avoid these issues, I consider only those youths who are living away from home on their interview date and who consider their current residence to be the same as their permanent residence. The data reveal that of these youths, 19 percent live alone, 41 percent live with a spouse or lover, and 16 percent live with siblings. Unfortunately, the presence of roommates cannot be explicitly identified.

These statistics indicate that although only one-fifth of youths live alone, less than one-quarter likely could be categorized as living with roommates in the sense of moving into an already established household with existing furniture and shared utilities. Leaving the parental home to live with a partner or a sibling likely requires a similar moving cost to living alone since youths would not be moving into already established households, which is an important difference between moving back with one's parents and moving in with contemporaries. However, the direct housing costs of living with others would possibly be lower than if living alone. The key implicit assumption in treating these living arrangements as "away" rather than as "home" is that the resource-sharing benefits and housing cost benefits for a youth are much larger when living with parents than when living with a sibling, partner, or roommate.

Parents Cannot Save

The model assumes that parents cannot save or borrow. This assumption is not innocuous but is necessary, for two reasons. First, and most important, the theoretical challenges to working with an imperfectly altruistic model without commitment in which both parties can save are overwhelming. Such models generally have a large set of Markov equilibria and to date are only understood in very stylized settings.[17] Second,

[17] See Barczyk and Kredler (2010*a*, 2010*b*) for a full discussion of the issues. They propose an equilibrium concept that is well defined. Their methods require continuous time and cannot be used in a setting with a discrete choice and fixed cost such as this one.

even if the theoretical challenges could be overcome (e.g., by making assumptions that remove strategic interaction between youths and parents), an additional continuous state variable would be introduced, and the computational complexity would render estimation infeasible.

The key implication of this assumption is that it forces financial transfers from parents to youths to lead to a reduction in parental consumption in the period that the transfer is made. If parents were able to hold assets, they could spread the consumption cost of transfers over future periods, making it cheaper in terms of lifetime utility. By ruling out parental saving, I am effectively limiting the extent to which parents will be willing to use financial transfers to offset the effects of labor market shocks to the youths. Relative to financial transfers, this makes coresidence a more effective form of transferring resources across generations. However, the alternative possibility (assuming that parents can save but youths cannot) would introduce an even larger bias into the value of coresidence as insurance: since youths would be unable to self-insure through savings, the value of any other potential insurance channel would be vastly overstated. For the purposes of the questions being addressed in this paper, it is thus crucial that youths have the possibility of self-insurance through savings.

Youths Do Not Contribute to Parental Expenses

The model assumes that youths do not pay housing costs (room and board) when living with parents or, equivalently, that parental transfers are nonnegative. Unfortunately, the NLSY97 does not ask questions about payments for room and board after the first wave (when youths are ages 13–17). However, the National Survey of Families and Households (NSFH), which is a cross-sectional survey about living situations, does contain evidence to justify this assumption.

I focus on the 2001 wave of the NSFH since it roughly corresponds to the sample period under consideration in the NLSY97. Table 5 shows weighted statistics on the fraction of 18–34-year-olds living with a parent who pay room and board and how much they pay.[18] For the full sample of youths, only around 21 percent pay anything, with a mean of less than $200. Focusing on the male subset under age 23 with a high school education or below (which is the sample most comparable to the

[18] The NSFH consists of three waves. For each respondent in the initial wave, a focal child, ages 18–34, was selected. In wave 3, each of these focal children who's primary place of residence was with the main respondent was asked about room and board. Because of the complex sampling design in the NSFH, it is not possible to construct a sample or weights so that statistics can be guaranteed to be representative. The weights used are those from the wave 1 respondents. Nonrandom attrition plus the fact that we look at focal children rather than respondents themselves means that even the weighted statistics may not be representative.

TABLE 5
STATISTICS ON ROOM AND BOARD FROM FOCAL CHILDREN IN NATIONAL SURVEY OF FAMILIES AND HOUSEHOLDS 2001

	All	Male	Male < Age 23	Male < Age 23 with ≤ High School Education
Pays for room and board (%)	21.3	20.6	12.3	17.9
Mean payment, conditional on paying ($)	196	179	159	143
Median payment, conditional on paying ($)	200	200	150	120
Individuals	338	182	106	50

NOTE.—Data come from wave 3 (2001) of the National Survey of Families and Households. Sample is all focal children who's primary place of residence is with the main respondent. Respondent weights from wave 1 are used in constructing all statistics.

NLSY97 sample), this fraction drops to 18 percent with a mean payment of $143. Hence, over 80 percent of youths do not make direct contributions when living with their parents. In addition, there are potentially large implicit transfers from parents when living at home (from meals, etc.), so that even for those youths who do contribute, the net transfer likely flows to the youth rather than to the parent.

IV. Estimation

A. *Estimation Strategy*

For certain parameters, structural estimation inside the model is less crucial than others. As such, the estimation approach involves fixing some parameters exogenously and estimating the remaining parameters using a set of moments from the NLSY97.

Externally Calibrated Parameters

Both parents and youths are assumed to have a risk aversion parameter, γ, equal to 1.5. The gross interest rate, R, is set equal to one because almost all wealth is held in the form of cash, checking accounts, or cars, all of which pay a zero (or perhaps negative) real financial return. The monthly unemployment benefit, b, is set at $500, and the monthly consumption floor is set at $100.[19] The distribution of parental income is

[19] Conditional on receiving unemployment benefits, the mean and median monthly benefits are $780 and $650. However, in the model all youths are eligible for unemployment benefits in all periods that they are not working. This is substantially more generous than in the US system, which requires that (i) a worker be laid off through no fault of his own, (ii) a worker satisfies an earnings or employment requirement over the previous year, and (iii) a worker collects unemployment benefits for no more than 26 weeks. To partially account for these differences, I reduce the mean amount in the data by around one-third to $500. Combining AFDC, food stamps, and WIC, the median monthly benefit

estimated from the NLSY97 data in a first stage. It is discretized to a four-point distribution, reflecting average parental income in each quartile.

The degree of economies of scale in the parental home, ϕ, and the costs of housing, χ, are particularly important parameters since they determine the economic benefits of coresidence. Although there are some data in the NLSY97 on rental costs, it is missing for much of the sample. The mean reported monthly rent for youths living away from home is \$380, based on 253 observations out of a total of 1,416. This number is significantly lower than what is suggested by the 2001 American Housing Survey. For renter-occupied units with low annual household income (\$11,700–\$17,550), the median monthly rent is \$601. In the model, χ refers to both direct and indirect costs of housing such as gas and electricity, so I set its value at \$650. The fixed cost of moving out of home, k, is set equal to 2 months of rent (\$1,300). At the estimated parameters, none of the allocations are significantly changed when housing costs are varied up or down by 15 percent or the fixed cost is set at either one or three times the monthly housing cost.

Unfortunately, there is not enough information in the NLSY97 data to identify the economies of scale in the parental home, ϕ, based on observed choices.[20] Since there are no consumption data in the NLSY97 and the Consumer Expenditure Survey (CEX) only measures expenditures at a household level, I rather calibrate ϕ on the basis of information from household-equivalence scales. I consider three of the most commonly used scales: the OECD (Organization for Economic Cooperation and Development) equivalence scale, the OECD modified scale, and the square-root scale. For each equivalence scale, I compute the percentage increase in income needed by a household to keep welfare constant when moving from a household with two adults to a household with three adults. These three scales give values of 41 percent, 33 percent, and 22 percent, respectively. Viewed through the lens of a static version of the model, these imply values for ϕ ranging from 0.20 to 0.42. I set $\phi = 0.4$, on the basis of these calculations.[21]

---

for the NLSY97 sample is \$220. Because eligibility and take-up are far from universal, I set the consumption floor at half this value. Results are not sensitive to other values in this range.

[20] With panel data on household-level expenditures, one could plausibly use changes in consumption expenditure when an additional member joins or leaves a household as a source of identification. See Lise and Seitz (2011) for an example of this approach using two-person husband-wife households. Identifying economies of scale within households is a long-standing research topic in applied econometrics, and one that is beyond the scope of this paper.

[21] Appendix H contains a detailed explanation of the mapping from equivalence scales to the weight on public goods in the utility function.

Internally Estimated Parameters

The approach for estimating the remaining 16 parameters is to use average moments over the age range from 17 to 22. I choose 21 moments that are sufficient to identify all the parameters, which I discuss in Section IV.B. The full set of moments is shown in table A2. Of these moments, the only one that relates labor market outcomes to coresidence outcomes is the cross-sectional difference in log earnings between youths living at home and away. This allows me to use the effect of labor market outcomes on coresidence dynamics as an informal out-of-sample test of the mechanisms at work in the model.

The estimated parameters are displayed in table 6 and include the labor market parameters (δ, λ_0^S, λ_0^R, μ_0^S, μ_0^R, σ_0, μ_d, μ_ϑ, σ_1); the preference shock parameters (α_z, β_z, σ_z^2, ρ_z); the altruism factor, η; the disutility of work, v; and the discount factor, β. These parameters are estimated using a simulated minimum-distance estimator with a diagonal weighting matrix.

B. *Model Fit, Parameter Estimates, and Identification*

Whenever structural parameters are estimated on the basis of simulated moments, a question of identification naturally arises. Although it is not possible to provide an analytic proof that the parameters are identified using a given set of moments, I address the question of identification in three ways. I start by following two approaches that are accepted as reasonable in the existing literature. First, I examined a numerical estimate of the Hessian of the minimum-distance criterion at the estimated parameter values and ensured that it is nonsingular. Second, I verified that the estimation strategy can recover good estimates of the structural parameters using data that are simulated from the model.[22] Third, below I provide an informal argument that each of the parameters has influence on a subset of the chosen moments and give some intuition for why this is the case. This approach should be persuasive since it delivers an understanding of why the available moments are sufficient to pin down the parameters. The fit of the model, as a function of age, is shown in figures 2 and 3. In this section, I discuss the parameter estimates in terms of their implication for the determinants of living arrangements and transfers.

[22] Both of these checks only suggest local identification. To check for other local minima, a thorough search of the parameter space was performed, and while other local minima were found, none of these improved on the fit of the estimated parameters.

TABLE 6
PARAMETER ESTIMATES

Parameter	Description	Estimate	90% Confidence Interval
Labor market:			
δ	Job destruction probability	.030	(.026, .039)
λ_0^S	*S*-type job arrival rate	.400	(.396, .515)
λ_0^R	*R*-type job arrival rate	.034	(.024, .038)
μ_0^S	Mean *S*-type log wage offer distribution	6.264	(6.056, 6.463)
μ_0^R	Mean *R*-type log wage offer distribution	7.613	(7.458, 8.252)
σ_0	SD log offer distribution	.943	(.612, 1.035)
μ_d	Mean *R*-type log wage growth	.052	(.041, .070)
μ_ϑ	Growth log experience effect	.011	(.008, .015)
σ_1	SD *R*-type log wage growth	.219	(.195, .311)
Preference shock:			
α_z	Intercept mean value of living away	−5.090	(−7.294, −3.946)
β_z	Age slope mean value of living away	.066	(.039, .104)
σ_z^2	Variance of value of living away	18.027	(14.191, 21.998)
ρ_z	Autocorrelation of value of living away	.974	(.951, .976)
Other:			
η	Altruism factor	.039	(.031, .042)
υ	Disutility of work ($\times 10^4$)	4.396	(.0003, 8.608)
β	Monthly discount factor	.964	(.959, .974)

NOTE.—Simulated minimum-distance estimator with diagonal weighting matrix. Bootstrap 90 percent confidence intervals in parentheses.

Labor Market Parameters and Moments

Since this is a search model, a standard identification challenge arises as a result of the fact that rejected job offers are not observed. In addition, the two types of jobs do not have directly observable analogues in the data. To achieve identification, I use a combination of functional form assumptions for the unconditional and the conditional wage offer distributions and the structural relationship between the disutility of labor and the two reservation wages. Eleven moments are used to pin down these nine labor market parameters. Although all moments are important for the values of all parameters, certain parameters have more influence on certain moments. The arrival rate of *S*-type jobs λ^S is identified from the probability of working, conditional on not working in the previous month, and the mean duration of unemployment. The job destruction rate (δ) is identified from the probability of not working, conditional on working in the previous month. Given the assumption of lognormality, the experience effect and the parameters of the wage offer distribution for *S*-type jobs ($\mu_0^S, \sigma_0, \mu_\vartheta$) are identified from the mean, variance, and mean growth in the distribution of log earnings, conditional on having not worked in the previous period. The same difficulty as in Flinn and Heckman (1982) applies here: it is the distri-

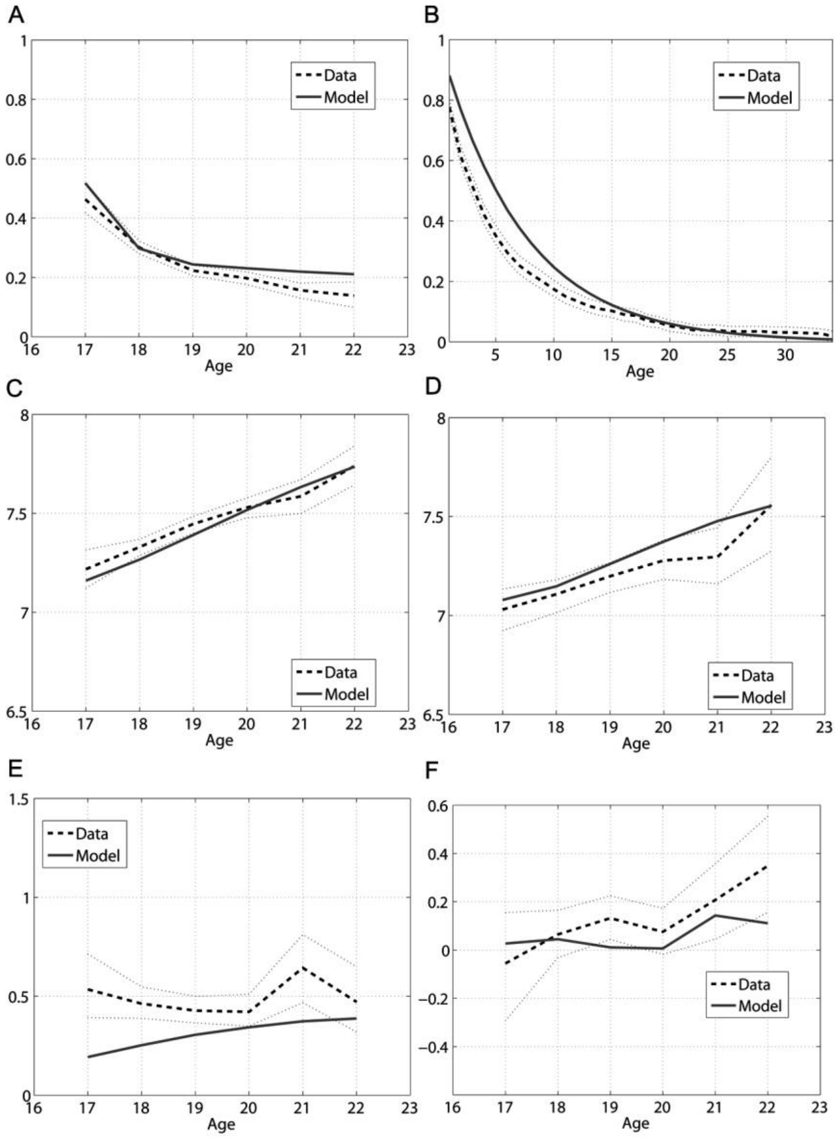

Fig. 2.—Model fit, labor market. Dashed line is data; dotted lines are 95 percent confidence interval for data; solid line is fit of model. *A*, Fraction not working; *B*, unemployment survival function; *C*, mean log earnings; *D*, mean log entry earnings; *E*, variance log earnings; *F*, away-home difference: mean log earnings.

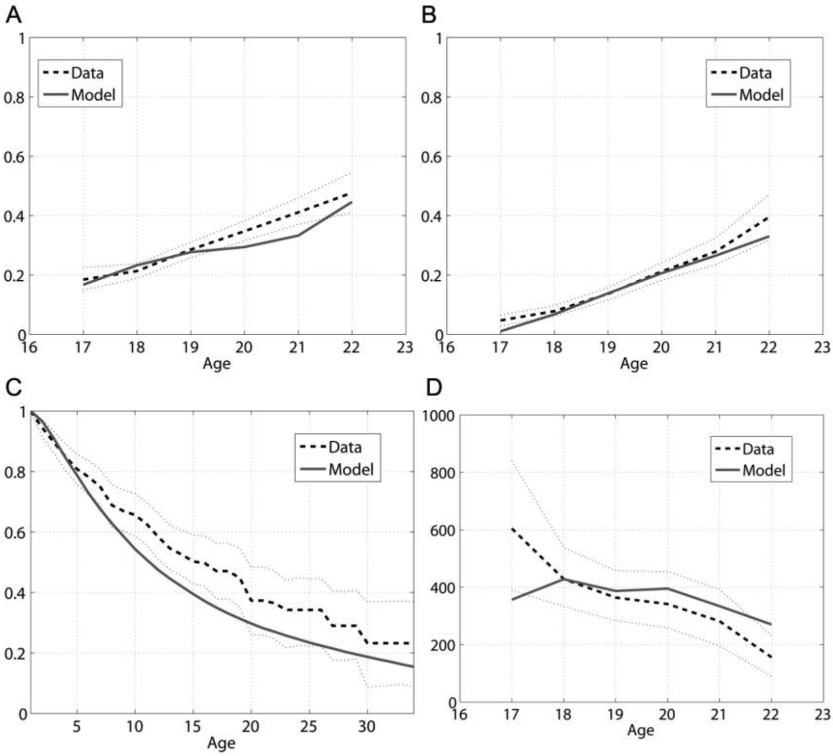

FIG. 3.—Model fit, coresidence. Dashed line is data; dotted lines are 95 percent confidence interval for data; solid line is fit of model. *A*, Fraction away from home; *B*, fraction ever moved back; *C*, duration back home survival function; *D*, average transfers.

butional assumptions alone that make it possible to distinguish a model with low arrival rates and a small mass of offers in the bottom of the wage distribution from one with a high arrival rate and more low-wage offers.

The mean and variance for earnings growth when employed in an *R*-type job (μ_d, σ_1) are pinned down by the mean and variance of the growth in log earnings, conditional on working in two consecutive periods, since in the model this is the only source of earnings growth relative to the growth in entry earnings. Relative to the parameters for *S*-type jobs, the arrival rate and mean wage offer for *R*-type jobs (λ_0^R, μ_0^R) are identified by second and higher moments of the overall wage distribution across both types (which have thicker tails than lognormal). I use the mean (in levels) of wages and variance of log wages.

Conditional on values for the labor market parameters, the disutility of work υ is identified by the average fraction of youths who are working in a given month. All of the key labor market characteristics are matched well by the model and are shown in figure 2.

Preference Shocks and Coresidence Patterns

The model is able to account for the key coresidence patterns in the data shown in figure 3. The intercept and slope in the mean utility from independence (α_z, β_z) are identified from the average fraction of youths living away from home and the growth in this fraction from age 17 to 22. The linear growth rate of preferences, β_z, is best measured in terms of the extent to which it accounts for this increase. When β_z is set to zero, with all other parameters left at their estimated values, the model generates 29 percent of the increase in the fraction of youths living away from home. This implies that just over two-thirds of the increase in the fraction living away is due to an increasing preference for independence. The other third is driven by purely economic factors: increasing earnings, employment, and the accumulation of assets.

The variance of preference shocks, σ_z^2, which determines the amount of heterogeneity in the relative preference for living away, is identified from the average difference in earnings between youths living at home and youths living away from home. Figure 2*F* shows that in the cross section, average earnings of working youths at home and away are similar, despite the evidence from Section II.A that earnings are important for coresidence transitions. If there were no preference heterogeneity ($\sigma_z^2 = 0$), then all coresidence movements would be driven by earnings and asset accumulation, and youths living away from home would, on average, have far higher earnings than youths living at home. As σ_z^2 increases, the amount of non–labor market heterogeneity increases. This additional heterogeneity reduces the cross-sectional away-home differential in earnings. As such, a fairly large amount of heterogeneity in preferences, σ_z^2, is needed to match the small difference in earnings between youths away and at home. Notice that the model is able to endogenously generate a small part of the gently increasing age profile of the away-home earnings difference. The reason is that at young ages, the only individuals who move out are those with a very strong preference for living away from home (high z). To finance their strong desire to live away from home, these youths lower their reservation wages and accept lower-paying jobs than youths who are living at home (and more often accept *S*-type jobs rather than wait for *R*-type jobs). This selection effect generates a small away-home earnings difference at the youngest ages. However, as the mean value of independence increases and youths have time to receive more offers and accumulate assets, the mix of youths who are living away from home shifts to comprise those who have received more favorable labor market shocks. The difference thus becomes larger at older ages. The high value for σ_z^2 ensures that there are always some low-earnings youths living away from home and some high-earnings youths living at home, as is implied by the small overall away-home earnings difference in the data.

The persistence of preferences, ρ_z, is identified by the within-person time series variation in parental coresidence. Three moments are used: (i) the monthly autocorrelation of coresidence outcomes, (ii) the fraction of youths who ever move back home at least once by age 22, and (iii) the mean duration of spells back home. To see why the last moment is important, consider an extreme environment in which there is no time series variation in z. In this case, the duration back home for a youth who moves back in response to an unemployment shock will be similar to the duration of unemployment. The extent to which spells back home are longer duration than spells out of work helps determine the extent to which movements back home were triggered by preference versus labor market shocks. The autocorrelation ρ_z is estimated to be 0.974 at a monthly frequency, which translates to an annual autocorrelation of 0.73. Thus, although there is a large amount of cross-sectional variation in the relative preference for living away from home (indicated by the high value of σ_z), there is much less within-person time series variation in preferences. This implies that although non–labor market heterogeneity plays a large role in explaining cross-sectional differences in coresidence outcomes, the labor market is the key factor in explaining individual movements in and out of the parental home.

To illustrate this point, I decompose coresidence patterns in the benchmark equilibrium. I do this by performing a standard within-groups/between-groups variance decomposition for cross-sectional (*a*) coresidence outcomes and (*b*) indicator variables for whether a youth moved in or out of home. These decompositions answer questions of the form: How much of the fact that one youth moved back home in a particular month, while another youth did not, is due to the fact that they received different histories of preference shocks?[23] I find that only 17 percent of

[23] Consider the cross-sectional variation in residence states, r_t. This can be decomposed as $\mathrm{Var}[r_t] = \mathrm{Var}[E(r_t|z^t)] + E[\mathrm{Var}(r_t|z^t)]$, where z^t denotes the entire history of preference shocks up to time t. The first term is the "between" component: variation in r_t that is due to cross-sectional differences in the history of realized preferences for living away from home. The second term is the "within" component: differences in coresidence states that exist even within groups of individuals who have experienced exactly the same history of preferences for independence. The fraction that is not accounted for by preferences, and hence is driven purely by labor market differences, is that due to the within component: $E[\mathrm{Var}(r_t|z^t)]/\mathrm{Var}[r_t]$. In order to calculate this fraction, it is necessary to calculate $E[r_t|z^t]$, which is a high-dimensional object. To do this, two approximations are made. First, the history is truncated after d periods. Second, I use a flexible nonparametric estimator for the conditional expectation, $E[r_t|z_t, \dots, z_{t-d}]$. Note that the decomposition could also have been specified in terms of histories of labor market shocks, w^t. The reason for preferring the decomposition in terms of preferences is that there is much less history

cross-sectional differences in coresidence are accounted for by labor market shocks. However, a far greater fraction of movements in and out of home are due to labor market events: 35 percent of movements back home and 72 percent of movements out of home.

An alternative way to assess the relative importance of the labor market and preferences for coresidence is to compare the benchmark model with a counterfactual environment where preference shocks are shut down but preference heterogeneity is retained. This is done by leaving all parameters at their estimated values, except for the transition matrix for preferences, which is replaced with an identity matrix. Hence, individuals do not experience changes to their value for z over time. I find that without shocks to preferences, the model generates 30 percent of the number of youths who ever move back home, suggesting that labor market shocks account for around one-third of spells back home by this measure.

Altruism Factor and Transfers

The altruism factor η is identified by the average level of financial transfers from parents to youths in a given year. As discussed in Section III.B, the optimal transfer decision is directly influenced by the weight that parents place on their children's utility. Figure 3*D* shows that the model does well in matching this average level and also the downward-sloping trend with age. The estimated value for η is 0.04.

This good fit for the average level of transfers masks cross-sectional variation in financial transfers that the model fails to capture well. In particular, the model generates a lower fraction of youths who receive a positive transfer compared with the data and a higher level of transfers conditional on receiving a positive transfer than in the data. The reason is due to the homotheticity of preferences that derives from altruism. This effect is easiest to see in the simpler static version of the model in Appendix C, where I show that the combination of altruism and identical preferences for youths and parents implies that transfers increase linearly with parental income. In the data, however, transfer amounts are relatively flat with parental income. To simultaneously match the overall level of transfers, the high fraction of youths receiving positive transfers, and flat transfers with parental income, the model would need to be changed in a way that could generate nonlinear Engel curves for par-

dependence in the effects of z than w, for the reasons discussed in the text. The calculation is done with $d = 3$, and the results are unchanged if a larger value is used.

ents' expenditures on youths. One such modification would be to allow for different curvatures of utility for the two generations.[24]

Discount Factor and Wealth

The discount factor β is identified from the mean level of assets at age 20. The extent to which assets are accumulated during these years, given the amount of risk that youths face and the implicit insurance from parental transfers and coresidence, identifies the degree of impatience. Expressed annually, the estimated discount factor is fairly low: 0.66, which reflects both (i) the relatively small degree of asset accumulation between ages 17 and 20 that is observed in the data and (ii) the propensity for many youths in the model to accept *S*-type jobs rather than waiting for *R*-type jobs, given the estimated differences in the long-term earnings growth that they generate.[25]

C. Effect of Labor Market on Coresidence Dynamics

The key economic forces in the model are those that relate the labor market to coresidence dynamics. Yet the only moment used in estimation that contains joint information on the labor market and coresidence is the cross-sectional difference in average log earnings between youths living at home and youths living away from home. This leaves open the possibility of using data on coresidence movements, conditional on labor market status, as an informal overidentification test of the mechanisms at work in the model. Table 7 shows the effect of employment status on the monthly probability of moving in and out of home in the data and in the estimated model. The model generates employment effects on the probabilities of moving that are of a similar magnitude but stronger than in the data.

[24] There are other simple changes to the model that would yield one extra parameter that could be used to match the fraction of youths receiving a positive transfer and the average level of transfers conditional on being positive. For example, one change would be a simple iceberg cost of transferring resources from the parent to the youth. However, such a mechanism would be somewhat artificial. Estimation of such a model would yield extremely altruistic parents (to match the high fraction of youths who receive transfers) and an extremely high cost of transferring resources (to match the low transfer amounts). An alternative modification to the model would be to allow for a negative correlation between altruism and parental income.

[25] Many other studies in a variety of contexts also suggest that males in this age group are relatively myopic. See Gruber (2000) for recent examples.

TABLE 7
MONTHLY PROBABILITY OF CORESIDENCE TRANSITIONS, BY LABOR MARKET STATUS (%)

	MOVE BACK HOME		MOVE OUT OF HOME	
	Data	Model	Data	Model
Overall	3.4	2.2	2.4	1.7
Working	3.1	1.7	2.7	2.3
Not working	4.9	4.2	1.9	.5
Difference	−1.8	−2.5	+.7	+1.8

V. Putting the Model to Work

A. *Coresidence as Insurance*

Measuring Insurance against Shocks

Consider a youth and a parent at the beginning of period t, after the realization of the shock (w_t, j_t, z_t) for that period. Recall that the state variables in the MPE are $x_t = (a_t, r_{t-1}, w_t, z_t, j_t, I)$, with a corresponding value function $Y_t(x_t)$ for the youth. I measure insurance as the degree to which a youth is indifferent between particular realizations of a shock.[26] Focusing on job destruction as a shock, define the difference in continuation values due to a job loss as

$$\Delta_t(x_t) \equiv Y_t(a_t, r_{t-1}, w_t, z_t, j_t, I) - Y_t(a_t, r_{t-1}, 0, z_t, j_t, I).$$

A youth is fully insured against a job loss if he is indifferent between losing and not losing his job; that is, if $\Delta_t(x_t) = 0$. When a youth is not fully insured, we can define the degree of partial insurance, $\xi_t(x_t)$, as the compensating asset variation that is necessary to make him indifferent between losing and not losing the job:

$$Y_t(a_t + \xi_t, r_{t-1}, 0, z_t, j_t, I) - Y_t(a_t, r_{t-1}, 0, z_t, j_t, I) = \Delta_t(x_t).$$

Thus, $\xi_t(x_t)$ is the answer to the question of how much additional wealth we would have to give a youth with state vector x_t to make him indifferent about becoming jobless.

Now consider a modification to the environment that removes a particular insurance channel. Denote the analogous continuation value difference in the resulting MPE as $\hat{\Delta}_t(x_t)$. Once again, define the extent of partial insurance, $\hat{\xi}_t(x_t)$, as the compensating asset variation for the job loss but valued according to the value functions in the benchmark equilibrium:

[26] This definition of insurance departs from some of the existing literature, by defining perfect insurance as a state of indifference rather than as equalization of the marginal utility of consumption. I refer to this alternative definition as *consumption* insurance and examine the consumption response to shocks in Sec. V.B.

$$Y_t(a_t + \hat{\xi}_t, r_{t-1}, 0, z_t, j_t, I) - Y_t(a_t, r_{t-1}, 0, z_t, j_t, I) = \hat{\Delta}_t(x_t).$$

The reason for using the value functions from the benchmark equilibrium is that we want to express the value differences in the benchmark and alternative environments using the same units. However, since the marginal value of assets may differ across the two equilibria, if we were to calculate $\hat{\xi}$ as the compensating asset variation implied by $\hat{Y}_t$, we would risk concluding that the utility loss from losing a job in one environment is larger than in another, simply because assets are not very valuable in that environment.[27] The way that $\hat{\xi}$ has been defined, it will always be the case that $\hat{\xi}_t(x_t) = \xi_t(x_t)$ whenever $\hat{\Delta}_t(x_t) = \Delta_t(x_t)$. We can then define the value of a particular insurance channel against the loss of a job at x_t as the proportionate increase in the cost of a job loss due to removing that channel:

$$\Omega_t(x_t) = \frac{\hat{\xi}_t(x_t)}{\xi_t(x_t)}.$$

There are at least two sets of benefits to measuring insurance in this way.

Substitution of independence for consumption.—Measuring insurance in terms of the smoothness of consumption is not appropriate (although I do this below, in order to connect with the existing literature, which has largely focused on this measure) since youths can adjust the inputs to their own welfare (consumption, labor supply, independence) in response to exogenous shocks. Consider, for example, a youth who moves back home as a result of a job loss shock. This youth is unambiguously worse off as a result of the shock (since he could always have quit his job), but his consumption may actually increase due to the public consumption in the parental home and reduction in housing costs. For this youth, the welfare cost of the job loss is not realized through a drop in consumption but rather through a loss of independence. However, if he were restricted from moving back home, he would retain his utility from independence but suffer a drop in consumption; $\Omega_t(x_t)$ takes both components of welfare into account, trading them off in the same way that the youth himself would.

Absence of level effects.—An arguably more standard way of measuring the welfare costs from the removal of insurance channels would be to simply compare the equilibrium value functions in the two environ-

[27] Consider, e.g., the case of removing unemployment insurance. Without unemployment insurance, assets are particularly valuable, so even if the utility loss from losing a job is large, the amount of assets that would be needed to compensate for the job loss may be small. Removing an insurance channel affects (i) the level of continuation values, (ii) the difference in continuation values between having and not having a job, and (iii) the marginal continuation value of additional assets. The measure of insurance defined in this section is designed to measure only the second of these effects.

ments. However, this comparison confounds two differences across the environment: differences in the overall level of welfare and differences in how welfare is affected by shocks. The insurance value of the change in the environment is only the latter effect, which is what $\Omega_t(x_t)$ measures. This distinction is particularly relevant in this model since the presence of public consumption inside the parental home and the existence of housing costs mean that the removal of the option of coresidence reduces the opportunity set for consumption. Even in a world without shocks, a youth would be worse off if he could not live at home, simply because there is less consumption available. The size of this effect is directly linked to the level of public goods in the parental home, determined by ϕ. However, in this paper I am interested in the second of these two effects: not the value of coresidence, per se, but the component of that value that is related to the effects of other shocks (e.g., job loss, wages) on welfare.

Welfare Cost of a Job Loss

I use the estimated model to calculate the value of different insurance channels for a typical 21-year-old male at different parts of the parental income distribution. I consider the welfare cost of a job loss for a youth with the mean preference for living away from home, the mean assets of this age group (approximately $1,000), and the mean earnings of this age group (approximately $2,000 per month). I do this separately for a youth who is working in an *S*-type and an *R*-type job.

Table 8 displays the results of these calculations for an *S*-type job in panel A and an *R*-type job in panel B. For each parental income group, the first two rows show the welfare cost of losing a job in the benchmark MPE. Expressed in terms of assets, the welfare costs of a job loss are around $3,400–$5,300 for an *S*-type job and are declining with parental income. These welfare costs translate to an immediate onetime transfer that is equivalent to between 2 and 3 months of earnings. This is consistent with an estimated arrival probability for *S*-type jobs of 0.4. For an *R*-type job, the welfare costs of a job loss with similar earnings are around three to five times larger and are increasing with parental income. The larger welfare costs reflect the higher future earnings growth that is lost when these jobs are lost, together with the fact that the arrival probability is around 10 times lower than for an *S*-type job.

The next three rows of each panel show how much these costs are increased by the removal of insurance channels. Removing the option to move back home increases the cost of a job loss substantially, but this effect decreases strongly as we move up the parental income distribution. The value of moving back home as insurance, as measured by Ω, is a factor of 20.9 (12.4) for a youth from the bottom quartile of the parental

TABLE 8
COST OF JOB LOSS AND VALUE OF INSURANCE

	Parental Income Distribution			
	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	A. Youth with *S*-Type Job			
Cost of job loss, ξ_t:				
Compensating asset transfer ($)	5,300	4,900	3,800	3,400
Number of months of earnings	2.9	2.7	2.1	1.9
Value of insurance channels, Ω_t:				
Option to move back home	20.9	33.1	4.7	3.1
Financial transfers	1.0	1.0	1.4	1.7
Unemployment benefits	2.6	2.5	3.7	4.1
	B. Youth with *R*-Type Job			
Cost of job loss, ξ_t:				
Compensating asset transfer ($)	16,000	16,700	17,800	18,400
Number of months of earnings	7.4	7.7	8.2	8.5
Value of insurance channels, Ω_t:				
Option to move back home	12.4	15.4	5.1	4.2
Financial transfers	1.0	1.0	1.2	1.3
Unemployment benefits	1.6	1.5	1.6	1.6

NOTE.—Values for a 21-year-old youth with median preference for living away from home and assets of $1,000; *S*-type job is for a youth with monthly earnings of $1,830, and *R*-type job is for a youth with monthly earnings of $2,170. Values of insurance (Ω_t) are expressed as multiple of cost of job loss (ξ_t).

income distribution with an *S*-type (*R*-type) job and decreases to 3.1 (4.2) for a youth from the top quartile. The main reason for these differences is that without the option for youths to move back home, wealthier parents compensate by increasing financial transfers when a youth becomes unemployed, while poorer parents cannot afford to do so. So the insurance value of being able to move back home is largest for those with the poorest parents, notwithstanding the fact that the level of public consumption provided in the parental home increases as we move up the parental income distribution.

Accordingly, the value of financial transfers moves in the opposite direction with parental income. Since youths with low parental incomes are less reliant on financial transfers for insurance in the benchmark equilibrium, removing this channel has only a small effect on the welfare cost of a job loss and virtually no effect for those in the bottom half of the distribution regardless of the type of job. However, for youths from wealthy families, the cost of removing financial transfers is larger: a factor of 1.7 (1.3) for a youth from the top quartile of the parental income distribution with an *S*-type (*R*-type) job. There are two contributing factors to the higher insurance value of financial transfers for these youths. First, restricting transfers has the direct effect of reducing their net income available for consumption. Second, restricting transfers forces some youths from wealthy families to move back home upon losing a job, reducing their utility from independence.

To put the size of these numbers into perspective, the last row of each panel in table 8 shows the value of the unemployment benefit as a channel of insurance in the model. Recall that unemployment is a $500 monthly benefit that is paid to all youths in any months that they are not working. Focusing on the loss of an *S*-type job (panel A), we see that for youths from the top half of the parental income distribution, the value of the option to move back home as an insurance channel is of a similar order of magnitude as the value of the unemployment benefit. However, for youths with poorer parents, the option to move back home is far more valuable than either financial transfers or unemployment benefits.

B. *Consumption Response to Shocks*

An alternative way to measure the value of insurance is to focus exclusively on consumption fluctuations and ignore compensating gains in utility from increased leisure (and reinforcing drops in utility from the loss of independence). Much of the existing literature on insurance against labor market shocks has followed this approach. In particular, Blundell et al. (2008) use data from the CEX and the Panel Study of Income Dynamics to measure the extent to which household consumption responds to household-level income shocks in the United States. Kaplan and Violante (2010) compare their findings with the predictions from a model in which the only mechanism for smoothing consumption is self-insurance through a risk-free security. They find that in such a world, consumption responds substantially more to income shocks than in the US economy, particularly for young households and for households that are borrowing constrained. In this section, I provide some indirect evidence that coresidence may constitute one of the additional mechanisms that young workers use to smooth consumption.[28]

Table 9 reports the average percentage drop in consumption associated with a job loss. In the benchmark equilibrium, this drop is 28 percent. When the option to move back home is removed, the consumption response from a job loss in the resulting equilibrium increases by 4 percentage points to 32 percent. When youths are restricted from living with their parents altogether, the consumption drop almost doubles to 53 percent. These findings suggests that coresidence has a sig-

[28] In this model, household formation is endogenous, and the focus is on individual-level consumption by the youth, whereas in the CEX data, households are taken to be a fixed unit, and consumption is measured at the household level. It is not possible to calculate corresponding individual-level consumption drops directly from CEX data.

TABLE 9
AVERAGE PERCENTAGE DROP IN CONSUMPTION IN RESPONSE TO LOSS OF A JOB

		Parental Income Distribution			
	Overall	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Baseline model:					
All	.28	.35	.26	.28	.24
High z_{it}	.53	.61	.58	.50	.45
Lower z_{it}	.23	.30	.20	.24	.20
Alternative models:					
No move back	.32	.39	.33	.29	.26
No coresidence	.53	.58	.58	.50	.44
No transfers	.23	.35	.26	.18	.14

nificant impact on the ability of youths to smooth consumption, through economies of scale and savings on direct housing costs.

However, the same cannot be said for financial transfers. When financial transfers are restricted, the consumption response to a job loss decreases by 5 percentage points. The reason is that removing financial transfers causes many youths to delay moving out of home, or to immediately move back home upon losing their job, since their parents cannot provide direct financial support. This means that relative to the benchmark, the utility drop from the loss of a job is realized through a loss of independence and not through a fall in consumption.

Table 9 also reports how these drops in consumption vary across the parental income distribution. The importance of the option of moving back home for consumption smoothing is most pronounced for youths from the poorest families. Restricting financial transfers has almost no effect on consumption drops for youths from the poorest families since these youths are not heavily reliant on financial transfers in the benchmark equilibrium. However, youths from the wealthiest families have large declines in the effect of a job loss on consumption when transfers are restricted since in the benchmark equilibrium they are heavily reliant on transfers to be able to live away from home.

A proxy for the option to move back home.—Structural estimation of the model parameters revealed both a high autocorrelation and a cross-sectional variance for the direct value of living away from home, z_{it}. This combination of parameters means that in the model there exist some simulated youths who have a very strong preference for a particular living situation, regardless of their labor market shocks. In particular, there are some simulated youths with a preference for living away from their parents that is so high that they will do so regardless of their labor market outcomes, at all ages. This feature of the model turns out to be useful for measuring the effect of the option to move back home. In the data, it is only possible to observe youths who actually do move back home, rather than youths who have a feasible option of moving back home should the need arise. However, in the model one can interpret a youth with a very high value for z_{it} as effectively not having an option to move back home. Since z_{it} is orthogonal to the realization of labor market shocks, these youths can be used as a control group for studying the effect of the option of living with one's parents on behavior.

I consider youths whose value of z_{it} is in the highest category throughout the sample period as the control group. The highest category is such that no youth in this category ever lives at home, regardless of other simulated state variables. In this sense, these are indeed youths without the option to live at home. Table 9 reports that consumption drops from a job loss separately for these youths (labeled high z_{it}) and all other youths (labeled low z_{it}). The results indicate that those youths without the option to move back home suffer consumption drops upon job loss that are on average 25–31 percentage points larger than for youths with this option, with the effect declining in the level of parental income.

C. Labor Market Outcomes

Implications for Labor Elasticities

Rios-Rull (1996), Gomme et al. (2004), Jaimovich and Siu (2009), and Dyrda et al. (2012) have all provided empirical evidence that the contribution of low-skilled young workers to aggregate labor market fluctuations is disproportionately large. In this section, I argue that recognizing that young people can have the option to live with their parents can shed light on this evidence and provides a partial explanation for these findings.

In models with a discrete individual labor supply decision, the aggregate labor supply elasticity is determined by the sensitivity of the distribution of reservation wages to changes in the distribution of wage offers (Chang and Kim 2007). In general, this elasticity is large whenever there is a large mass of workers with reservation wages near their current offered wage; that is, a large number of workers are close to being indifferent about working at their offered wage.

When youths do not have the option to live at home, fewer young workers are marginal: a higher fraction of all job offers are accepted, leaving less room for labor supply to respond to changes in the distribution of wage offers. When youths have this option, those who choose to live at home have access to public consumption, even when they are not working. This results in a larger number of offers being rejected: a greater fraction of young workers are nearly indifferent about accepting job offers.

TABLE 10
IMPLIED LABOR ELASTICITY IN MODEL: JOB ACCEPTANCE PROBABILITIES

		PARENTAL INCOME DISTRIBUTION			
	OVERALL	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Baseline model:					
All	.26	.25	.25	.26	.27
Home ($r_{it-1} = 0$)	.24	.24	.24	.24	.25
Away ($r_{it-1} = 1$)	.32	.32	.31	.31	.33
High z_{it}	.31	.30	.30	.30	.33
Lower z_{it}	.25	.25	.24	.25	.26
Alternative models:					
No coresidence	.30	.29	.29	.30	.33
No transfers	.25	.25	.25	.26	.25

To illustrate the potential size of these effects, table 10 reports the elasticity of labor supply with respect to a job offer, ε^o, which is defined as the probability of accepting an offer, conditional on receiving one while unemployed: $\varepsilon^o \equiv \Pr(h_t = 1 | w_t > 0, w_{t-1} = 0)$. The overall probability that an offer is accepted in the baseline model is 0.26. However, there is a large difference between youths living at home and youths living away from home: the acceptance rate for youths living at home, 0.24, is 25 percent lower than for youths living away from home, 0.32. A corollary is that when coresidence is ruled out, the overall acceptance rate increases substantially to 0.30. Exploiting the differences in behavior between youths with low values versus high values of z_{it}, as in the previous section, yields a similar finding, a difference in acceptance rates of around 20 percent.

In contrast, when financial transfers are restricted but coresidence is allowed, the probability of accepting an offer slightly declines, relative to the benchmark. This is because without financial transfers, a number of youths from wealthy families delay departure from the parental home.

Implications for Earnings

In addition to affecting short-term labor supply decisions, the option to live at home can have important implications for longer-term labor market outcomes. There are two channels through which this happens in the model. First, as discussed above, the option to live at home leads youths to raise their reservation wages and wait longer for better initial wage draws. Since there is a substantial degree of persistence in wages in both the model and the data, this leads to higher earnings at later ages. Second, the option to live at home makes it more likely that a youth will work in an *R*-type job, which are the jobs that generate wage growth. This happens because (i) offers for *R*-type jobs arrive less fre-

TABLE 11
Labor Market Outcomes at Age 23

	Overall	Parental Income Distribution			
		Quartile 1	Quartile 2	Quartile 3	Quartile 4
	A. Fraction in *R*-Type Job (%)				
Baseline model:					
All	22.1	21.8	23.1	20.4	23.4
High z_{it}	19.7	18.2	23.2	14.1	23.4
Lower z_{it}	22.6	22.5	23.1	21.6	23.3
Alternative models:					
No coresidence	20.6	18.7	20.4	19.4	24.0
No transfers	21.7	21.8	23.1	20.0	21.8
	B. Median Monthly Earnings ($)				
Baseline model:					
All	2,040	1,960	2,070	2,060	2,055
High z_{it}	1,990	1,850	1,960	2,040	2,140
Lower z_{it}	2,050	1,990	2,100	2,070	2,040
Alternative models:					
No coresidence	1,940	1,830	1,915	2,010	2,040
No transfers	1,930	1,870	1,960	1,980	1,930

quently than *S*-type jobs, so on average a longer period of unemployment must be suffered before such an opportunity arises, and (ii) *R*-type jobs have riskier and delayed payoffs, both of which are more palatable to youths who have the option of living with their parents.

Table 11 reports results on the quantitative magnitude of these effects from the estimated model. Panel A shows the fraction of working youths who hold an *R*-type job at age 23. Overall, this fraction is 22.1 percent. When the option to live with one's parents is removed, this fraction drops to 20.6 percent. This means that without coresidence there would be around 7 percent fewer youths working in riskier, high-growth jobs, for the reasons discussed above. Using preferences for living in the parental home, z_{it}, as a proxy for the option to move back home yields slightly larger effects: 22.6 versus 19.7 percent, which is an increase of around 15 percent. Panel B of table 11 shows that these differences in job acceptance decisions manifest as higher wages for youths with the option to live at home. When the possibility of coresidence is allowed, youths have wages that are on average 5 percent higher than when youths are precluded from living with their parents.

D. Savings Behavior

In this section, I illustrate how the Samaritan's dilemma (Bruce and Waldman 1990) implicit in the model of altruism and financial transfers

TABLE 12
Average Asset Holdings at Age 23, Relative to Baseline Model

	Overall	Parental Income Distribution			
		Quartile 1	Quartile 2	Quartile 3	Quartile 4
Average assets at age 23 ($)	6,550	6,825	6,900	6,100	6,530
Change from baseline (%):					
No move back	16	18	22	14	12
No coresidence	1	0	6	2	−3
No transfers	10	3	3	17	17

distorts youths' savings behavior. The effect is reminiscent of the findings in Hubbard et al. (1995), who showed that the asset-based means testing in the US social insurance system provides a disincentive to accumulate private savings and can thus help explain the low observed savings rates of households with low lifetime earnings. For young people, the implicit conditioning of parental transfers on youths' assets has an analogous effect.

These effects are illustrated in table 12, which reports average asset holdings at age 23 in the benchmark model and the percentage difference in the counterfactual economies with each dimension of parental support removed. When the option to move back home is removed, average asset holdings increase by 16 percent. Without this option, youths have a stronger incentive to accumulate wealth for precautionary reasons. Similarly, when financial transfers are removed, average assets increase by 10 percent. Youths in the bottom half of the parental income distribution are more reliant on the option to move back as a form of insurance than on direct financial assistance; hence, their savings are more distorted by the former channel than the latter. Youths in the top half of the parental income distribution rely more heavily on financial transfers, so the reverse is true for this group. Note that coresidence plays an important role in these counterfactual savings outcomes, even in the experiment in which financial transfers are restricted. In this economy, coresidence rates increase, particularly at younger ages, which provides youths with the means to accumulate the precautionary savings that will compensate for the lack of future transfers.

Removing the possibility of coresidence entirely has two offsetting effects on youth savings. On the one hand, without the option to live at home there is a stronger precautionary motive for savings. On the other hand, not being able to live at home leads to a large cost in terms of current resources, from both rent obligations and the lack of public goods in the parental home. This makes it more difficult to save. On average, the two effects approximately cancel out, and the effect on overall savings is around 1 percent.

VI. Conclusions

In this paper I have argued that the option for young people to live with their parents is a valuable channel of insurance against labor market risk, with important implications for behavior. Being able to live with their parents allows young people to search longer for jobs that have better prospects for future earnings growth. By reducing the consumption response to labor market shocks, the option to live at home can help explain why young households appear to have access to insurance possibilities over and above that implied by self-insurance through savings. Through its effect on the incentives to accumulate precautionary savings, parental support can generate a plausible mechanism for explaining the low savings rates of low-skilled youth. By raising the reservation wages of young people, the option to live at home can help explain the high aggregate labor elasticity of young workers.

At a policy level, the implications of parental coresidence are potentially far-reaching. The fact that living arrangements respond endogenously to the realization of labor market shocks suggests the possibility of substantial crowding out by social insurance programs. Since many public programs are designed to insure against the same types of idiosyncratic labor market shocks that living arrangements respond to, it is important to consider the impact on coresidence when evaluating their welfare implications. Examples of policies that condition on living arrangements include those that require means testing at a household level (e.g., food stamps) and those that link benefit entitlement to the structure of households (e.g., Temporary Assistance for Needy Families).

Appendix A

Data and Sample Selection

A. National Longitudinal Survey of Youth 1997

The sample is drawn from the NLSY97, which is a longitudinal survey of 8,984 individuals from the cohort born between 1980 and 1984. The survey contains extensive information on labor market behavior and educational outcomes, together with detailed information on the youth's family and community background. Interviews have been conducted approximately annually since 1997.

Coresidence Variables

In principle, information about parental coresidence in the NLSY97 may be obtained in two ways. The simplest way is through the household roster, which records the relationship to the youth of all individuals living in the household at the time of the interview. However, using the household roster as the basis for coresidence information has two problems. The first is that it provides only an (approximately) annual snapshot of living arrangements. This means that using the household roster does not allow one to observe the circumstances that surround any change in living arrangements. The second problem with using the household roster is that it does not necessarily correspond to the current residence of the youth. Rather, it refers to the residence that the youth considers to be his primary residence. For example, a youth who has recently moved out of home may report not living with either of his parents but still report that his parents' home is his primary residence. The rosters are thus not reliable indicators of current residential status, particularly for youths whose coresidence status has recently changed.

Instead, I focus on the set of retrospective questions about monthly coresidence that were asked in rounds 2–6 (1998–2002). At each interview, these questions asked respondents to list each period of 1 month or more in which they lived separately from each of their parents. A parent is defined in the NLSY97 as a biological, step, adoptive, or foster parent.

This specific wording of the retrospective monthly coresidence questions differed slightly across waves. The main differences are that in the first five rounds, the respondent is given a list of each calendar month since the previous interview and asked about coresidence status in each of the months. However, in round 6 the youth is instead asked to report directly the calendar months in which changes in coresidence status took place. For each parent figure, the youth was asked the following set of questions:

Rounds 2–5 (*1998–2001*).—"Now I'm going to ask you about your parents and any other people you consider to be parent figures. I will ask about each parent separately. My computer tells me that at the time of your last interview, you were living [under joint custody/blank] with your [mother (figure)/father (figure)][name]. Was this information correct on [date of last interview] when we had your last interview?" Questions are then asked about updating any changes in the parent's characteristics from the previous wage:

> Are you currently living with [him/her] full-time, living with [him/her] as part of a joint custody arrangement, or not living with [him/her] at all? Since [date of last interview], has there been a continuous period of one month or more when you and your [mother (figure)/father (figure)] lived in different places? If you were temporarily away at summer camp, but lived with your [mother (figure)/father (figure)] before and after that time, please include those months as months you were living with [him/her].
>
> Since [date of last interview] what months have you lived with your [mother (figure)/father (figure)] at least some of the time? If you were temporarily away at summer camp or on vacation, but lived with your [mother (figure)/father (figure)] before and after that time, please include those months as months you were living with [him/her].
>
> Now I'd like to ask you about parents and parent figures you weren't living with at the time of our last interview. Since [date of last interview], has there been a continuous period of one month or more when you and your [mother (figure)/father (figure)] lived in different places? Since [date of last interview] what months have you lived with your [mother (figure)/father (figure)] at least some of the time? If you were temporarily away at

> summer camp or on vacation, but lived with your [mother (figure)/father (figure)] before and after that time, please include those months as months you were living with [him/her].

Round 6 (*2002*).—"Since [date of last interview], has there been a continuous period of one month or more when you and your [mother (figure)/father (figure)] lived in different places?" "Since [respar1datefill] what month did you [resparent1firstnext] stop living with your [mother (figure)/father (figure)]? What month did you [resparent1firstnext] start living with your [mother (figure)/father (figure)]?" "Was there another period of a month or more when you did not live with your [mother (figure)/father (figure)]?" This group of questions is then repeated in a loop.

By piecing together the responses to these coresidence questions across rounds, it is possible to reconstruct a monthly panel of parental coresidence outcomes for each respondent. A youth is defined as living away from his parents in a given month only if he is observed to be not living with any of his living parent figures for the entire month. Conversely, a youth is defined to be living at home if he reports living with at least one parent figure at any point during the month. This implies that only spells away from home that are longer than 1 month's duration are considered to be valid spells in the analysis, and all spells back home will be recorded as lasting at least 1 month.

Labor Market Variables

Labor market variables in the NLSY97 are constructed from three sources: (i) the employer roster, which records details about each job that the youth has held; (ii) the employment event history, which is a weekly record of which employers the youth worked for in a given week; and (iii) created variables for total hours worked in the week and hourly compensation for each job in each week. The challenge is to construct monthly variables for employment, hours, and earnings from these weekly data. To do this, I define a week as falling in a particular month if the start date was on or before the twenty-eighth of the month (twenty-fifth for February). This means that each month has either 4 or 5 assigned weeks.

A youth is defined as working in a particular month if he or she is recorded as working for at least one employer during at least 1 week in that month. Monthly earnings for working youths are defined as 52/12 times average weekly hours in that month, multiplied by average hourly compensation, where the averages are taken across all jobs and all working weeks in the month. Where hours are missing but wages are available and the youth reports working full-time, 40 hours are assumed. Some youths report unrealistically high wages and hours. I deal with this misreporting by setting weekly hours from above 100 to 100 and hourly wages from above $75 to $75, in 2007 dollars. None of the results are sensitive to the choice of these thresholds.

Asset Data

From round 4 onward, information about assets and debts was only collected from youths in the first interview after they turn 20 or 25. This means that for the relevant ages in the sample, asset information at age 20 is the most useful. However, in the first three rounds, the NLSY97 collected information on all independent youths. Having finished school, living alone, or being 18 are each sufficient to be considered independent. Thus, in the first three waves, asset data are available for all youths in the sample. My measure of assets includes all financial assets and vehicles less financial debts and money owed in respect to vehicles owned. Financial assets include businesses, pension and retirement accounts, savings accounts, checking accounts, stocks, and bonds. There is some top coding in the data: the top 2 percent in each category is set at the average of that group.

Transfers Data

In rounds 1–7, youths also reported parental transfers in the past year. Aside from allowances and parental loans, respondents stated the amount of money they received from each parent or guardian. Example questions for parental transfers include, "Other than allowances, did your (mother/female guardian) give [you/you or your spouse/partner] any money during 2002? Please include any gifts in the form of cash or a check but do not include any loans from your (mother/female guardian)." And, "How much money did your (mother/female guardian) give [you/you and your spouse/partner] during 2002?"

The transfers data are top coded for the top 2 percent of the sample. In my baseline sample, 32 percent of male youths report receiving a positive amount of financial transfers in the past year. Among those reporting positive transfers, the mean (median) amount received was $1,120 ($450).

Unfortunately, information about negative transfers (payments from youths to their parents) was only collected in round 1. In the first round, respondents were asked whether they made regular payments to their parents during the previous year. For respondents living at home, this survey collected information on money paid for room and board. These questions were discontinued from round 2 onward.

Sample Selection

The baseline sample restricts attention to youths who are never observed to participate in any type of postsecondary education (sample A). As discussed in the main text, to address the potential concerns regarding endogeneity of the education decision, I also compare the results from this sample with those from two alternative samples that implement the restriction to low-skilled youth in different ways. Sample B selects on the basis of low test scores, which are a strong predictor of future college participation. This is done by retaining only youths who scored in the bottom quartile of the combined Armed Services Vocational Aptitude Battery, as constructed by NLS staff. The tests were administered before the first wave, during the NLSY97 screening process. Hence, selection in this

TABLE A1
Number of Respondents Lost at Each Stage of Sample Selection

	Number Lost	Number Remaining
Raw National Longitudinal Survey of Youth 1997		4,599
Drop respondents with 1998 interview missing	316	4,283
Drop respondents still in high school	740	3,543
Drop respondents observed in postsecondary education	1,683	1,860
Drop respondents with gaps in coresidence data	16	1,844
Drop respondents who ever have all parents dead	17	1,827
Drop respondents ever in military	131	1,696
Drop respondents ever in jail	124	1,572
Drop respondents ever in a shelter for homeless/abused	3	1,569
Drop months at age 17	9	1,560
Drop respondents with gaps in employment data	69	1,491
Final sample	1,491	
Month-youth observations	36,222	
Mean observations per person	24.3	
Median observations	21	

sample is based on a purely exogenous variable, and there are no issues of endogeneity of education. Sample C is less restrictive than the baseline sample. Rather than dropping any youth who is ever observed to participate in postsecondary education of any type, I only drop youths who are "traditional" college participants—those youths who start college immediately after graduating from high school or within 1 year of graduating. The purpose of sample C is to retain youths who may attend college part-time.

The NLSY97 is itself not representative, due to an oversampling of black and Hispanic youths, as well as nonrandom attrition. Hence, in all estimations I use a custom set of cross-sectional weights to account for oversampling and attrition for 1997–2002. The weights are based on the characteristics of youths who are present in all six rounds (1997–2002). See http://www.nlsinfo.org/web-investigator/custom_weights.php for information on the construction of customized weights for use with the NLSY97 (see table A1).

Data Moments for Estimation

Table A2 reports the 21 moments used in estimation for the three samples.

B. Current Population Survey

I use the Basic Monthly Surveys from the CPS from 1979 to 2010. I adopt the CPS definition of a household: all persons who occupy a dwelling unit. A dwelling unit is defined as "a room or group of rooms intended for occupation as separate living quarters and having either a separate entrance or complete cooking facilities for the exclusive use of occupants." This definition largely coincides with the notion of a household from the point of view of economic theory. An individual is considered to be living with his or her parents if his or her rela-

TABLE A2
MOMENTS USED IN ESTIMATION

	Sample A	Sample B	Sample C
Labor market moments:			
Mean earnings	2,221	2,224	1,999
Mean log earnings	7.462	7.464	7.348
Variance log earnings	.486	.493	.516
Mean log entry earnings	7.187	7.191	7.096
Variance log entry earnings	.534	.581	.507
Growth mean log earnings	.104	.096	.064
Growth mean log entry earnings	.106	.104	.067
Variance growth log earnings	.508	.531	.613
Fraction not working	.244	.237	.307
Probability stop work	.043	.042	.049
Probability start work	.154	.153	.121
Mean unemployment duration	5.287	5.322	6.834
Coresidence moments:			
Fraction living away	.298	.298	.265
Growth in fraction living away	.058	.059	.048
Fraction ever moved back	.396	.386	.347
Autocorrelation coresidence	.933	.931	.932
Away-home log earnings difference	.164	.146	.153
Growth away-home log earnings difference	.034	.026	.076
Mean duration back home	17.5	17.5	19.0
Other moments:			
Mean transfers, including zeros	357	392	521
Mean assets age 20	3,664	4,597	4,031

tionship to the household head is that of child. The data are de-trended with a Hodrick-Prescott filter, allowing for a structural break in January 1991 to account for updating the census-based population weights that induced a discrete change in the number of different types of individuals.

Appendix B

Details of Econometric Model for Duration Analysis

Consider a setting where time, $t \in (0, \infty)$, is continuous and a youth who is currently living away from home has individual characteristics (possibly time varying) given by X_t. Denote the hazard of moving back home by $\theta(t, X_t)$. I assume a proportional hazards representation, where the log hazard is separable between a baseline time-varying hazard, $\theta_0(t)$, and a linear function of X_t:

$$\theta(t, X) = \theta_0(t) e^{\beta' X_t}.$$

The survivor function is thus given by

$$S(t, X) = \exp\left\{-e^{\beta' X} \int_0^t \theta_0(s)\, ds\right\}.$$

Since the data are discrete and organized in months, we can only observe that a youth moved back during a particular month. It is thus useful to derive the discrete time hazard of moving back home in a given month. Define month t as the month between time $t-1$ and t. Then the hazard during month t is

$$h_t(X) = 1 - \exp\left\{-e^{\beta' X}\int_{t-1}^{t}\theta_0(s)\,ds\right\}.$$

If we define $\gamma_t = \log[\int_{t-1}^{t}\theta_0(s)\,ds]$ as a measure of duration dependence (it reflects the integrated baseline hazard during month t), then we can write the discrete time hazard as

$$\log\{-\log[1 - h_t(X)]\} = \beta' X_t + \gamma_t,$$

which is a complementary log-log regression. I estimate the parameters β and report e^{β}. This is the multiplicative effect on the underlying continuous time baseline hazard: it tells how much the baseline hazard of moving back home is affected at all durations, given a change in the variable X_t.

For the model of moving back home, duration dependence through γ_t is proxied with a polynomial in the age of the youth. This is because it is likely that it is the age of the youth rather than the number of months since he or she first moved out that is important for moving back home. This also avoids the issue of left censoring for moving back home since for a number of spells we do not observe when the youth first moved out of home.

The model can also allow for unobserved heterogeneity in the form of random effects:

$$\log\{-\log[1 - h_t(X)]\} = \beta' X_t + \gamma_t + \alpha_i.$$

I estimated models with two types of unobserved heterogeneity: one in which α_i is distributed according to a gamma distribution and one in which α_i is distributed according to a normal distribution. Both models delivered quantitatively similar results, so to conserve space, only results from the normal model are reported. In these models, α_i reflects a spell-level, rather than a person-level, fixed effect. This specification allows different spells away from home for the same youth to be characterized by different unobserved characteristics.

Appendix C

A Static Version of the Game

I describe the structure and Nash equilibria of a static version of the game in the full model with exogenous labor supply. Since this version of the model admits closed-form solutions, it is useful to demonstrate some of the key mechanisms at work in the full model.

Consider a static version of the game in which youth income is exogenous, there is no fixed cost of moving out, there are no savings, $\phi = 1$ (full economies of scale), and $\gamma = 1$. In this simplified version of the game, the only actions are the residential choice $r \in \{0, 1\}$ for the youth and the transfer amount $T \in [0, I^p]$ for the parent. The wage offer for the youth, w; the income of the parent, I^p; and the utility from independence are taken as exogenously given parameters. The payoffs in this game are given by

$$\text{Youth:} \quad U^y = \log[g^y + (1-r)g^p] + rz,$$

$$\text{Parent:} \quad U^p = \log(g^p) + \eta U^y,$$

with the following budget constraints:

$$g^y + r\chi \leq w + T,$$

$$g^p + T \leq I^p,$$

$$T \geq 0.$$

Note that when the youth lives at home, it is always optimal for the parent to set $g^p = I^p$. Hence, the payoffs can be written as functions of (r, T):

$$\text{Youth:} \quad U^y(r, T) = \log[w + T - r\chi + (1-r)I^p] + rz,$$

$$\text{Parent:} \quad U^p(r, T) = \log(I^p - T) + \eta U^y(r, T),$$

with $T = 0$ when $r = 0$. The assumption of log utility is not essential to obtain a closed-form solution; however, it simplifies the algebra.

A. *Best Response of Parent*

The optimal transfer for a parent with a youth living away from home is given by

$$T^*(1) = \frac{1}{\eta + 1}[\eta I^p - w + \chi].$$

I assume that $w \leq \eta I^p + \chi$ to simplify the algebra—this assumption just says that parental income plus housing costs is much larger than the youth's earnings. Note that this implies that $T^*(1) \geq 0$.

B. *Best Response of Youth*

For a given transfer amount T, the youth will live away from home if $U^y(1, T) \geq U^y(0, T)$, where

$$U^y(1, T) = \log(w + T - \chi) + z,$$

$$U^y(0, 0) = \log(w + I^p),$$

which generates a reservation transfer for living away from home given by

$$\tilde{T} = \frac{w(1 - e^z) + I^p + e^z\chi}{e^z},$$

where I have assumed $w \geq \chi$, so that it is always feasible for the youth to live away from home, and $z > 0$. The best response of a youth is to live away from home if $T \geq \tilde{T}$ and at home otherwise.

C. *Nash Equilibria*

There is a Nash equilibrium where the youth lives away from home as long as $T^*(1) \geq \tilde{T}$ and a Nash equilibrium where the youth lives at home if $\tilde{T} \geq 0$. Since $T^*(1) \geq 0$, a Nash equilibrium will always exist; however, it need not be unique.

The latter Nash equilibria (where the youth lives at home) are sustained by the parent offering any transfer $T<\tilde{T}$ if the youth is to move out. This implies that if $T^*(1)\geq\tilde{T}$ there will be both an away equilibrium and a home equilibrium.

The payoffs for the youth and the parent in each equilibrium are given by

$$U^y(0,0) = \log(w+I^p),$$

$$U^y(1,T^*(1)) = \log\left[\frac{\eta}{\eta+1}(w-\chi+I^p)\right]+z$$

and

$$U^p(0,0) = \log(I^p)+\eta\log(w+\phi I^p),$$

$$U^p(1,T^*(1)) = \log\left[\frac{1}{\eta+1}(w-\chi+I^p)\right]+\eta\log\left[\frac{\eta}{\eta+1}(w-\chi+I^p)\right]+\eta z,$$

respectively.

The youth prefers the away equilibrium whenever

$$z\geq\log\left(\frac{w+I^p}{w-\chi+I^p}\right)-\log\left(\frac{\eta}{\eta+1}\right),$$

but the parent prefers the home equilibrium if

$$z\leq\log\left(\frac{w+I^p}{w-\chi+I^p}\right)-\log\left(\frac{\eta}{\eta+1}\right)+\frac{1}{\eta}\log\left[\frac{(\eta+1)I^p}{w-\chi+I^p}\right].$$

If $z\geq\log[(w+I^p)/(w-\chi+I^p)]-\log[\eta/(\eta+1)]+(1/\eta)\log[(\eta+1)I^p/(w-\chi+I^p)]$, then both the youth and the parent prefer the away equilibrium; that is, it Pareto dominates the home equilibrium. If $z<\log[(w+I^p)/(w-\chi+I^p)]-\log[\eta/(\eta+1)]$, then both the youth and parent prefer the home equilibrium: it Pareto dominates. In both of these cases, if a timing protocol were specified, then the subgame perfect equilibrium would be the same, regardless of who chose first. However, if z lies in the interval between these two thresholds, there is disagreement as to which equilibrium is preferred. The youth prefers the equilibrium in which he lives away, and the parent prefers the equilibrium in which the youth lives at home. If the youth were to choose first, then he would choose to live away, and this would be the subgame perfect equilibrium. However, if the parent were to choose first, he or she would choose $T^*=0$, and the home equilibrium would ensue. In this case, the timing protocol does matter.

Appendix D

Markov-Perfect Equilibrium

The MPE of the game can be described by a set of Bellman equations. Define $Y_t^m(x_t^m)$ and $P_t^m(x_t^m)$ as the expected discounted value along the equilibrium path at the beginning of phase m of the period t stage game, for the youth and the parent, respectively. The four phases of the stage game, and the corresponding state variable, x_t^m, are outlined in table D1. Optimal decisions for the youth and the parent are denoted with an asterisk (*). The value functions for the youth are given in equations (D1)–(D3). Equation (D1) is the expected value at the beginning of period t, before the current period shocks have been realized; (D2) describes the discrete residence decision, taking into account the equilibrium transfer strategy of the parent; and (D3) is the labor supply and savings decision, which takes into account future values along the equilibrium path, summarized in $Y^1_{t+1}(x^1_{t+1})$:

$$Y^1(x^1_t) = \sum_{w_t, j_t, z_t} Y^2_t(x^2_t) \Pr(w_t, j_t | w_{t-1}, j_{t-1}, h_{t-1}) \Pr(z_t | z_{t-1}), \tag{D1}$$

$$Y^2_t(x^2_t) = \max_{r_t} \{Y^4_t(a_t, r_{t-1}, r_t, w_t, j_t, z_t, T^*_t(x^3_t), g^{p*}_t(x^3_t))\}, \tag{D2}$$

$$Y^4_t(x^4_t) = \max_{h_t, a_{t+1}} \{U^y_t + \beta Y^1_{t+1}(x^1_{t+1})\}, \tag{D3}$$

subject to (5), where $U^y_t = \{[c^{y(1-\phi)}_t G^\phi]^{1-\gamma}/(1-\gamma)\} - h_t v + r_t z_t$, and $G = g^y_t + (1 - r_t) g^p_t$.

Equations (D4) and (D5) describe the problem faced by a parent along the equilibrium path. Equation (D4) is the expected value for the parent at the beginning of period t, which depends on the residence choice of the youth. Equation (D5) is the optimal transfer and public good decision, which takes into account the induced labor supply and savings decisions of the youth as well as the residence decision in the following period.

$$P^1_t(x^1_t) = \sum_{w_t, j_t, z_t} P^3_t(a_t, r_{t-1}, r^*_t(x^2_t), w_t, j_t, z_t) \Pr(w_t, j_t | w_{t-1}, j_{t-1}, h_{t-1}) \Pr(z_t | z_{t-1}), \tag{D4}$$

$$P^3(x^3_t) = \max_{c^p_t, g^p_t, T_t} \{U^p_t + \eta U^{y*}_t + \beta P^1_{t+1}(a^*_{t+1}(x^4_t), r_t, h^*_t(x^4_t), w_t, j_t, z_t)\}, \tag{D5}$$

subject to (6), where $U^p_t = [c^{p(1-\phi)}_t g^{p\phi}_t]^{1-\gamma}/(1-\gamma)$, $U^{y*}_t = \{[c^{y*}_t(x^4_t)^{(1-\phi)} G^{*\phi}]^{1-\gamma}/(1-\gamma)\} - h^*_t(x^4_t) v + r_t z_t$, and $G^* = g^{y*}_t(x^4_t) + (1 - r_t) g^p_t$.

The optimal transfers for a youth living away from home are given by

$$\tilde{\phi}(c^p_{jt} + g^p_{jt})^{-\gamma} \le \eta \frac{\partial Y^4_t}{\partial a_t} + \beta \frac{\partial \tilde{P}^1_{t+1}}{\partial a_t}, \tag{D6}$$

$$T_{jt} = 0 \text{ if inequality is strict,} \tag{D7}$$

where $\tilde{\phi} = [(1-\phi)^{1-\phi}\phi^\phi]/(1-\gamma)$ and $\tilde{P}^1_t \equiv P^1_t(x^1_t) - \eta Y^1_t(x^1_t)$ is the parents' direct value function, that is, the portion of their continuation value along the equilibrium path that is due to their own consumption only and not their altruistic feelings. The left side of equation (D6) is the marginal value of current period consumption for the parent in which the optimal split of resources between the

TABLE D1
State Variables for Different Phases of the Stage Game

Phase	Conditioning Variable	Choice	By Whom	Strategies	Value Function
1	$x^1_t \equiv a_t, r_{t-1}, h_{t-1}, w_{t-1}, j_{t-1}, z_{t-1}$	w_t, j_t, z_t	Nature		$Y^1_t(x^1_t), P^1_t(x^1_t)$
2	$x^2_t \equiv a_t, r_{t-1}, w_t, j_t, z_t$	r_t	Youth	$r^*_t(x^2_t)$	$Y^2_t(x^2_t)$
3	$x^3_t \equiv a_t, r_{t-1}, r_t, w_t, j_t, z_t$	T_t, g^p_t	Parent	$T^*_t(x^3_t), g^{p*}_t(x^3_t)$	$P^3_t(x^3_t)$
4	$x^4_t \equiv a_t, r_{t-1}, r_t, w_t, j_t, z_t, T_t, g^p_t$	h_t, a_{t+1}	Youth	$h^*_t(x^4_t), a^*_{t+1}(x^4_t)$	$Y^4_t(x^4_t)$

two types of consumption goods has been imposed. Note that when the youth lives away from home, the decision about how to divide spending between the two types of goods is not affected by altruism, so the optimal decision is to allocate an exogenous fraction of spending to each good. The first term on the right side of equation (D6) is the marginal value to the parent of an additional unit of resources for the youth, before the youth's current period labor supply and savings decisions have been made. This marginal value is the altruism factor multiplied by the marginal value of resources to the youth himself. The second term on the right side of (D6) is the discounted direct value to the parent of entering the next period with the youth holding an extra dollar of assets. This value reflects the lower expected future transfers that would eventuate along the equilibrium path if the youth had an extra unit of resources.

Appendix E

Pareto Efficiency

There are a number of reasons why the MPE allocations may be inefficient, relative to an environment where parents and youths can commit at $t = 0$ to fully history-dependent allocations. First, since parents cannot commit to transfers before youths make their coresidence decision, there may be inefficient delays in moving out of home and inefficient movements back home. Second, since youths cannot commit to accept a job before parents make their transfer decisions, there may be inefficiently low transfers. Finally, a version of the Samaritan's dilemma is at work, whereby youth's savings are inefficiently low because they seek to raise their marginal value of resources in order to induce higher transfers from parents.

To examine how severe are these inefficiencies, and thus the sensitivity of the results to alternative choices about how to determine allocations, I construct the Pareto-efficient frontier between parents and youths, at the estimated parameter values. I then look at the difference in welfare and allocations between the MPE and nearby points on the Pareto frontier. For a given value for the youth, $\overline{V_0^y}$, define efficient allocations to be the subset of S that satisfy

$$\overline{V_0^p} = \max_{s \in S} P_0^1 \qquad \text{subject to } Y_0^1 \geq \overline{V_0^y} \tag{E1}$$

and the efficient frontier as the subset of the locus $(\overline{V_0^y}, \overline{V_0^p})$ for which the constraint $Y_0^1 \geq \overline{V_0^y}$ binds. Due to the presence of altruism, there may be feasible values for $\overline{V_0^y}$ for which this constraint does not bind. In these cases, both the youth and the parent can be made better off by increasing the welfare of the youth. Clearly such allocations are not Pareto efficient and so are not included as part of the Pareto frontier.

Note that the problem in (E1) can be rewritten as

$$\max_{s \in S} \tilde{P}_0^1 + (\eta + \lambda) Y_0^1, \tag{E2}$$

where λ is a Lagrange multiplier that can be interpreted as a relative Pareto weight on the youth. Since only the combined altruism factor/Pareto weight, $\eta + \lambda$, is important for determining allocations, the assumption of efficiency alone is not sufficient for identification of η. I construct the Pareto frontier by fixing η at its estimated value and solving the problem in (E2) for different values of λ.

A. *Game Is Nearly Efficient*

In Kaplan (2010), I report computed efficient frontiers $(\overline{V_0^y}, \overline{V_0^p})$ for the four parental income groups and the corresponding value pairs (Y_0^1, P_0^1) from the MPE. It turns out that the actual inefficiencies in the game are extremely small since (Y_0^1, P_0^1) lie very close to the efficient frontier. In the cases in which the game is almost efficient, the MPE lies close to a point on the Pareto frontier that puts essentially no direct weight on the utility of the youth (i.e., $\lambda \approx 0$). If the game were indeed to generate efficient allocations, then this would necessarily be the case. To see why, it is useful to compare equation (D6), which determines transfers for youths living away from home in the game, with the corresponding equation in the efficient allocations:

$$\tilde{\phi}(c_{jt}^p + c_{jt}^g)^{-\gamma} \leq (\eta + \lambda)\frac{\partial Y_t^4}{\partial a_t} + \beta\frac{\partial \tilde{P}_{t+1}^1}{\partial a_t}, \tag{E3}$$

$$T_{jt} = 0 \text{ if inequality is strict.}$$

It is clear that if the value functions in (D6) and (E3) are to coincide, then the only way that the game could generate an efficient level of transfers is if $\lambda = 0$.

B. *Why Prefer the Game as the Baseline?*

If the inefficiencies generated by the game are so small (and the game is more difficult to compute), then why focus on the MPE, which requires additional assumptions about timing and commitment, as the preferred model of behavior? There are at least four reasons. First, the game is intuitive and generates some outcomes that appeal to introspection about the nature of parent-youth interactions. For example, parents may make substantial transfers, even though they would prefer the youth to live at home. Also, parents cannot control the labor supply and savings of nonresident youths directly but can only partially influence them through their choice of financial transfers. If a youth has a strong enough preference for independence, he will move out regardless of the parent's actions.

Second, the particular specification of the game implicitly assumes that parents and youths cannot commit to future decisions. This seems more in touch with reality than the assumption implicit in the Pareto-efficient allocations—that parents and youths can commit at age 17 to a full set of contingent allocation rules for coresidence, labor supply, consumption, and savings.

Third, there is an important advantage of the game in terms of identification of structural parameters. Under the assumption of Pareto efficiency, only the combined Pareto-weight and altruism factor, $\eta + \lambda$, is identified, which makes it difficult to use the model to do policy experiments and examine counterfactual exercises. Only a locus of possible counterfactual outcomes is identified, indexed by how the estimated value for $\eta + \lambda$ is split between η and λ. The implicit bargaining weight, λ, may respond endogenously to changes in the environment, an issue that is not of concern when using the game to determine allocations. If the efficient allocations were implemented through some decentralized system, then λ would reflect the implicit bargaining power given to the youth and hence the resulting point on the Pareto frontier. However, depending on the details of the decentralization, changes in the environment may change the effective value of λ, leading to a different point on the Pareto frontier being chosen. In the dynamic game, the assumption about the timing of actions pins down the effective bargaining power of youths and parents. Finally, it is useful to be able to connect with the existing literature on parent-youth interactions, which has predominantly used a noncooperative game-theoretic approach to model behavior.

Appendix F

Tax Function

I use a tax function that includes three types of tax: payroll, federal, and state. Payroll taxes are composed of two parts: (i) social security tax of 6.2 percent of annual income up to $102,000 and (ii) a Medicare levy of 1.45 percent of annual income with no limit. For federal income taxes, I define net income as gross income minus a standard deduction of $5,350 and a personal exemption of $3,400. I then use the progressive tax rates for a single with no dependents for 2007 on the basis of this net income. I assume that state income taxes are 2.5 percent of gross income less a deduction for federal taxes plus another $2,500. All calculations are based on annual income, by multiplying the monthly income by 12 and dividing the resulting tax bill by 12.

Appendix G

Numerical Solution of the Model

The model is solved by backward induction from the terminal value functions that are described in Section III. The distribution of preference shocks, z_{it}, is discretized to a five-point stationary Markov chain with second moments defined by ρ_z and σ_z^2. The grids are then transformed to yield a distribution with mean $\alpha_z + (t-1)\beta_z$. Value functions and decision rules are solved on a grid with 16 points for assets and seven points for public consumption inside the parental home. Linear interpolation (bilinear interpolation for two-dimensional problems) is used to evaluate values between grid points. The asset decision for the youth is solved using a golden search with multiple starting values at each point in the state space. The discrete choices (coresidence, labor supply) are solved by interpolating the choice-specific value functions at the relevant stage of the game.

Minimization of the simulated minimum-distance objective function is performed using a trust-region method for nonsmooth nonlinear least squares functions developed by Zhang, Conn, and Scheinberg (2010). The algorithm works by forming a sequence of smooth approximation to each of the moment conditions being targeted. I use multiple restarts from many different starting points in the parameter space to ensure that a global maximum is found. Confidence intervals are obtained by bootstrap.

Appendix H

Calibrating ϕ from Equivalence Scales

Consider the problem of a one-person household with Cobb-Douglas utility over two consumption goods, c and g, and income I:

$$\max_{c,g} g^{\phi} c^{1-\phi},$$

subject to

$$g + c = I.$$

Such a household sets $g^* = \phi I$ and $c^* = (1-\phi)I$ and obtains an indirect utility function $V^1(I) = \phi^{\phi}(1-\phi)^{1-\phi}y$.

Now consider a two-person household (whose members I will denote as p and y) with total income I. With a unitary model for decision making in the household, and equal weights on the two members, the household solves the following problem:

$$\max_{c,g} g^{\phi} c_y^{1-\phi} + g^{\phi} c_p^{1-\phi},$$

subject to

$$g + c_y + c_p = I.$$

The solution to this problem is to set $g = [(1+\phi)/2]I$ and $c_y = c_p = [(1-\phi)/2]I$. The indirect utility function for each of the two household members is $V^2(I) = [(1+\phi)/2]^{\phi}[(1-\phi)/2]^{1-\phi}I$.

The value of ϕ that is consistent with an equivalence scale e is defined as the ϕ such that $V^1(I) = V^2(Ie)$. This gives

$$e(\phi) = 2\left(\frac{\phi}{1+\phi}\right)^{\phi}.$$

Note that $e(0) = 2$, so that with no economies of scale, income needs to be doubled when moving from a one-person to a two-person household in order to keep welfare constant. Similarly, $e(1) = 1$, so that with full economies of scale, no additional income is required to keep welfare constant when a second member is added to the household.

I consider the additional income required when adding a second adult member to a two-adult household that is implied by three commonly used equivalence scales. I focus on the addition of a second adult, rather than a child, because 17–23-year-olds are better thought of as adults rather than children for consumption purposes. The OECD equivalence scale implies $e = 1.41$, the OECD modified scale implies $e = 1.33$, and the square-root scale implies $e = 1.22$. Using the formula above that relates e to ϕ, these numbers imply a value for ϕ between 0.20 and 0.42. I choose 0.4, which lies in this range.

References

Acemoglu, D., and R. Shimer. 1999. "Efficient Unemployment Insurance." *J.P.E.* 107 (5): 893–928.

Altonji, J., F. Hayashi, and L. Kotlikoff. 1997. "Parental Altruism and Inter Vivos Transfers: Theory and Evidence." *J.P.E.* 105 (6): 1121–66.

Attanasio, O., and S. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *J.P.E.* 104 (6): 1227–62.

Barczyk, D., and M. Kredler. 2010*a*. "Altruistically-Motivated Transfers under Uncertainty." Working paper, Universidad Carlos III de Madrid.

———. 2010*b*. "A Dynamic Model of Altruistically-Motivated Transfers." Working paper, Universidad Carlos III de Madrid.

Becker, G. 1974. "A Theory of Social Interactions." *J.P.E.* 82 (6): 1063–93.

Becker, S., S. Bentolila, A. Fernandes, and A. Ichino. 2010. "Youth Emancipation and Perceived Job Insecurity of Parents and Children." *J. Population Econ.* 23 (3): 1047–71.

Bethencourt, C., and J. Rios-Rull. 2009. "On the Living Arrangements of Elderly Widows." *Internat. Econ. Rev.* 50 (3): 773–801.

Bianchi, S., V. Hotz, K. McGarry, and J. Seltzer. 2006. "Intergenerational Ties: Alternative Theories, Empirical Findings and Trends, and Remaining Challenges." In *Caring and Exchange within and across Generations*, edited by A. Booth, N. Crouter, S. Bianchi, and J. Seltzer. Washington, DC: Urban Inst.

Blundell, R., L. Pistaferri, and I. Preston. 2008. "Consumption Inequality and Partial Insurance." *A.E.R.* 98 (5): 1887–1921.

Bruce, N., and M. Waldman. 1990. "The Rotten-Kid Theorem Meets the Samaritan's Dilemma." *Q.J.E.* 105 (1): 155–65.

Buck, N., and J. Scott. 1993. "She's Leaving Home: But Why? An Analysis of Young People Leaving the Parental Home." *J. Marriage and Family* 55 (4): 863–74.

Card, D., and T. Lemieux. 1997. "Adapting to Circumstances: The Evolution of Work, School, and Living Arrangements among North American Youth." Technical report, NBER, Cambridge, MA.

Chang, Y., and S. Kim. 2007. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations." *A.E.R.* 97 (5): 1939–56.

Costa, D. 1999. "A House of Her Own: Old Age Assistance and the Living Arrangements of Older Nonmarried Women." *J. Public Econ.* 72 (1): 39–59.

Danforth, J. 1979. "On the Role of Consumption and Decreasing Absolute Risk Aversion in the Theory of Job Search." In *Studies in the Economics of Search*, vol. 123, edited by S. A. Lippman and J. J. McCall, 109–31. New York: North-Holland.

Da Vanzo, J., and F. Goldscheider. 1990. "Coming Home Again: Returns to the Parental Home of Young Adults." *Population Studies* 44 (2): 241–55.

Deaton, A., and C. Paxson. 1994. "Intertemporal Choice and Inequality." *J.P.E.* 102 (3): 437–67.

Dyrda, S., G. Kaplan, and J. Rios-Rull. 2012. "Business Cycles and Household Formation: The Micro vs. the Macro Labor Elasticity." Technical report, NBER, Cambridge, MA.

Ermisch, J. 1999. "Prices, Parents, and Young People's Household Formation." *J. Urban Econ.* 45 (1): 47–71.

Ermisch, J., and P. Di Salvo. 1997. "The Economic Determinants of Young People's Household Formation." *Economica* 64 (256): 627–44.

Fernandez-Villaverde, J., and D. Krueger. 2002. "Consumption and Saving over the Life Cycle: How Important Are Consumer Durables?" In *Proceedings of the*

2002 North American Summer Meetings of the Econometric Society, edited by D. K. Levine and W. Zame, 1–55. New York: Econometric Soc.

Flinn, C., and J. Heckman. 1982. "New Methods for Analyzing Structural Models of Labor Force Dynamics." *J. Econometrics* 18 (1): 115–68.

Fogli, A. 2004. "Endogenous Labor Market Rigidities and Family Ties." Working paper, Fed. Reserve Bank Minneapolis.

Goldscheider, F., and C. Goldscheider. 1999. *The Changing Transition to Adulthood: Leaving and Returning Home.* Vol. 17. Thousand Oaks, CA: Sage.

Gomme, P., R. Rogerson, P. Rupert, and R. Wright. 2004. "The Business Cycle and the Life Cycle." *NBER Macroeconomics Ann.*, 415–61.

Gruber, J. 2000. "Risky Behavior among Youths: An Economic Analysis." Technical report, NBER, Cambridge, MA.

Hayashi, F., J. Altonji, and L. Kotlikoff. 1996. "Risk-Sharing between and within Families." *Econometrica* 64 (2): 261–94.

Hubbard, R., J. Skinner, and S. Zeldes. 1995. "Precautionary Saving and Social Insurance." *J.P.E.* 103 (2): 360–99.

Jaimovich, N., and H. Siu. 2009. "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility." *A.E.R.* 99 (3): 804–26.

Kaplan, G. 2010. "Moving Back Home: Insurance against Labor Market Risk." Staff Report 449, Fed. Reserve Bank Minneapolis.

———. 2011. "Inequality and the Lifecycle." Working Paper no. 11-014, Univ. Pennsylvania, Penn Inst. Econ. Res.

Kaplan, G., and G. Violante. 2010. "How Much Consumption Insurance beyond Self-Insurance?" *American Econ. J.: Macroeconomics* 2 (4): 53–87.

Lentz, R., and T. Tranaes. 2005. "Job Search and Savings: Wealth Effects and Duration Dependence." *J. Labor Econ.* 23 (3): 467–89.

Lise, J. 2011. "On-the-Job Search and Precautionary Savings: Theory and Empirics of Earnings and Wealth Inequality." Working Paper W11/16, Inst. Fiscal Studies, London.

Lise, J., and S. Seitz. 2011. "Consumption Inequality and Intra-household Allocations." *Rev. Econ. Studies* 78 (1): 328–55.

Livshits, I., J. MacGee, and M. Tertilt. 2007. "Consumer Bankruptcy: A Fresh Start." *A.E.R.* 97 (1): 402–18.

Low, H. 2005. "Self-Insurance in a Life-Cycle Model of Labour Supply and Savings." *Rev. Econ. Dynamics* 8 (4): 945–75.

Low, H., C. Meghir, and L. Pistaferri. 2010. "Employment Risk and Wage Risk over the Life Cycle." *A.E.R.* 100 (4): 1432–67.

Manacorda, M., and E. Moretti. 2006. "Why Do Most Italian Youths Live with Their Parents? Intergenerational Transfers and Household Structure." *J. European Econ. Assoc.* 4 (4): 800–829.

McElroy, M. 1985. "The Joint Determination of Household Membership and Market Work: The Case of Young Men." *J. Labor Econ.* 3 (3): 293–316.

Pezzin, L., R. Pollak, and B. Schone. 2007. "Efficiency in Family Bargaining: Living Arrangements and Caregiving Decisions of Adult Children and Disabled Elderly Parents." *CESifo Econ. Studies* 53 (1): 69–96.

Rios-Rull, J. 1996. "Life-Cycle Economies and Aggregate Fluctuations." *Rev. Econ. Studies* 63 (3): 465–89.

Rosenzweig, M., and K. Wolpin. 1993. "Intergenerational Support and the Life-Cycle Incomes of Young Men and Their Parents: Human Capital Investments, Coresidence, and Intergenerational Financial Transfers." *J. Labor Econ.* 11 (1): 84–112.

———. 1994. “Parental and Public Transfers to Young Women and Their Children.” *A.E.R.* 84 (5): 1195–1212.

Sakudo, M. 2007. “Strategic Interactions between Parents and Daughters: Coresidence, Marriage, and Intergenerational Transfers in Japan.” Manuscript, Univ. Pennsylvania.

Storesletten, K., C. Telmer, and A. Yaron. 2004. “Consumption and Risk Sharing over the Life Cycle.” *J. Monetary Econ.* 51 (3): 609–33.

Zhang, H., A. Conn, and K. Scheinberg. 2010. “A Derivative-Free Algorithm for Least-Squares Minimization.” *SIAM J. Optimization* 20 (6): 3555–76.