Author's personal copy

Journal of Financial Economics 96 (2010) 331–344

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Limited participation and consumption-saving puzzles: A simple explanation and the role of insurance ☆

Todd Gormley [a], Hong Liu [b,c], Guofu Zhou [b,c,*]

[a] The Wharton School, University of Pennsylvania, USA
[b] Olin Business School, Washington University in St. Louis, USA
[c] China Academy of Financial Research, China

ARTICLE INFO

Article history:
Received 2 April 2009
Received in revised form
29 May 2009
Accepted 8 June 2009
Available online 18 January 2010

JEL classification:
D11
D91
G11
C61

Keywords:
Limited participation
Insurance
Saving
Consumption

ABSTRACT

In this paper, we show that the existence of a large, negative wealth shock and insufficient insurance against such a shock could explain both the limited stock market participation puzzle and the low-consumption–high-savings puzzle. We then conduct an empirical analysis on the relation between household portfolio choices and access to private insurance and various types of government safety nets. The empirical results demonstrate that a lack of insurance against large, negative wealth shocks is positively correlated with lower participation rates and higher saving rates. Overall, the evidence suggests an important role of insurance in household investment and savings decisions.
© 2010 Elsevier B.V. All rights reserved.

1. Introduction

As Campbell (2006) states: "Textbook financial theory implies that all households, no matter how risk averse, should hold some equities if the equity premium is positive." But, it is well documented that less than 50% of households in the US, and far less in many other countries, participate in the stock market, and that this limited participation has a significant impact on the equilibrium risk premium, diversification discount, liquidity, and market crashes.[1] It is also well documented that, relative to their income, households in developing countries consume less and save more than those in developed countries (Cao and Modigliani, 2004). This gap is puzzling especially given the explicit effort by governments in many developing countries to discourage saving and promote consumption as a way to reduce dependence on foreign demand for economic growth.

☆ We thank Thorsten Beck, Jian Cai, John Campbell, Phil Dybvig, Michael Faulkender, Campbell Harvey, Juhani Linnainmaa, Mark Loewenstein, Anthony Lynch (American Finance Association discussant), David McLean (Financial Intermediation Research Society discussant), G. William Schwert (managing editor), Costis Skiadas, Rong Wang, Kangzhen Xie, conference participants at the 2009 American Finance Association meetings and 2008 Financial Intermediation Research Society meetings, seminar participants at Washington University in St. Louis, and especially Miles Kimball (referee) for his insightful and detailed comments that substantially improved the paper.

* Corresponding author at: Olin Business School, Washington University in St. Louis, Campus Box 1133, One Brookings Drive, St. Louis, MO 63130, USA. Tel.: +1 314 935 6384.
E-mail address: zhou@wustl.edu (G. Zhou).

0304-405X/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jfineco.2010.01.001

[1] See, for example, Mankiw and Zeldes (1991), Haliassos and Bertaut (1995), Heaton and Lucas (2000), Guiso, Haliassos, and Jappelli (2003), Basak and Cuoco (1998) and Huang and Wang (2007).

Author's personal copy

While several possible explanations exist for the limited participation puzzle and the low-consumption–high-savings (LCHS) puzzle, there is little overlap in these explanations. The existing explanations for limited participation rates generally rely on either transactions costs (Vissing-Jorgensen, 2002), the existence of both learning and short-sale constraints for investors whose income has a positive correlation with the stock market (Linnainmaa, 2005) or nonstandard preferences, in which nonparticipation can be optimal for investors with ambiguity or disappointment aversion (Dow and Werlang, 1992; Ang, Bekaert, and Liu, 2005; Epstein and Schneider, 2007; Cao, Wang, and Zhang, 2005). Existing explanations for the LCHS puzzle, however, generally take different routes and tend to rely on demographic structure and inflation (Cao and Modigliani, 2004), credit constraints and a desire to purchase durable goods (Chamon and Prasad, 2006), or a combination of habit formation and rapid economic growth (Carroll and Weil, 1994; Carroll, Overland, and Weil, 1997, 2000).

In this paper, we provide a simple model to show that a single driving force, insufficient insurance against large negative shocks, could help explain both the limited participation and the LHCS puzzles. We then conduct an extensive empirical analysis, at both the country and household level, to analyze the relation between stock market participation, saving, and insurance. We find substantial evidence that a lack of insurance against large, negative wealth shocks is strongly correlated with both lower participation rates and higher saving rates. The evidence strongly supports the theoretical result of our model and Elmendorf and Kimball (2000) that insurance could play a role in both household investment and savings decisions.

In our model, to help illustrate the potential role of insurance in resolving both puzzles, we highlight how exposure to a large, negative wealth shock, such as loss of a job or a health emergency, affects household portfolio choices. We show that if there exists a negative wealth shock and an investor needs almost all current wealth to finance subsistence-level consumption, then a rational investor saves almost all of his current wealth and does not participate in the stock market at all, even if the probability of such a shock is arbitrarily small and the expected stock return is high.[2] This is because the investor needs almost all of his current wealth to absorb the negative wealth shock and to maintain the subsistence level of consumption if such a shock occurs. This explanation is economically intuitive and does not require nonstandard preferences, fixed costs, or a positive income beta with short-sale constraints. It also does not require income growth or demographic shifts.

Our model implies an important role insurance markets can play in participation and saving decisions, which has not received much attention in the asset pricing literature. If there is an insurance market for hedging the wealth shock, an investor optimally participates in the stock market, consumes more, and saves less. If the probability of the negative wealth shock is small and thus buying insurance is inexpensive, then the investor could even consume almost all of his current wealth and save virtually nothing if his expected future wealth is high. The presence of government safety nets, which insure against such adverse shocks, has similar implications. Overall, the model indicates that availability of insurance can dramatically change an investor's consumption, saving, and investment decisions.[3]

Country-level empirical analysis provides strong support for a connection between insurance, participation rates, and saving. We find that stock market participation rates across countries are positively correlated with the presence of a large, private insurance market. This correlation holds strongly even after we control for other country-level characteristics shown to be important in the existing literature, including measures of investor protection, trust, legal origin, and economic development. In addition, access to insurance markets is negatively correlated with saving rates of households, particularly in developing countries where households face arguably greater exposure to large, negative wealth shocks. The presence of public insurance, as captured by government safety nets such as a large social security program, is also positively related to participation levels and negatively related to saving rates. Finally, we find that, consistent with the model prediction, participation and saving rates are negatively correlated at the country level.

Our empirical evidence is not limited to country-level correlations between the availability of insurance, participation rates, and saving rates. We also find evidence that differences in the generosity of unemployment insurance programs across US states is positively related to stock market participation levels in those states. This finding holds even if we account for individual household characteristics, such as age, income, wealth, race, marriage, risk tolerance, and after controlling for state-level differences in the average level of education and economic development. Consistent with our theory, the positive relation between unemployment insurance generosity and participation levels is also particularly strong in those US states where a larger fraction of workers are exposed to employment risk, as captured by greater employment in declining industries. This evidence suggests that social safety nets and access to insurance affect individuals' portfolio decisions even in developed countries.

Finally, we find evidence that the likelihood of a US household holding risky stocks is positively correlated with the extent to which the household is insured against an adverse health shock. US households that are insured against large, adverse health shocks—as captured by having an employer-sponsored health policy, life insurance policy, or long-term care policy—are more likely to

[2] This is qualitatively consistent with Polkovnichenko (2007), who finds that, with uninsurable stochastic income and habit-formation preferences, investment becomes more conservative.

[3] This theoretical result also conforms with recent policy advice given by US Treasury Secretary Timothy F. Geithner. News reports indicate that Secretary Geithner has urged China to stimulate domestic consumption by improving its social safety net through better pensions and health care. See Wall Street Journal Asia (2009).

Author's personal copy

hold risky stocks after controlling for various household characteristics including age, wealth, education, and income.[4]

Overall, our paper finds strong empirical evidence that access to both public and private insurance is related to households' portfolio decisions. The evidence suggests that access to insurance could play an important role in explaining both the limited participation puzzle and the LCHS puzzle and in understanding household investment behavior and asset pricing in general.

Our empirical evidence also provides substantial support to the large literature on precautionary saving against background risk (e.g., Kimball, 1992, 1993; Eeckhoudt and Kimball, 1992; Elmendorf and Kimball, 2000; Carroll, 2001; Carroll and Kimball, 2008). Kimball (1993) demonstrates that with the standard risk aversion, in the presence of the zero-mean background risk, an investor invests less in the risky asset and saves more. Elmendorf and Kimball (2000) discuss the effect of undesirable risk on both consumption and risky asset holding. Although they focus on the effect of income taxes in providing implicit insurance, the math clearly applies to other types of insurance, and they show that large shocks can have powerful effects. Our model extends Elmendorf and Kimball (2000) by linking background risk to limited participation and showing that, no matter how small this background risk is (in the sense that the probability of such a negative wealth shock can be arbitrarily small), it can prevent an investor from investing any amount of cash into risky assets.

The rest of the paper is organized as follows. In Section 2, we use a simple model to illustrate the impact insurance could have on household participation and saving rates. In Section 3, we show the empirical relation between household portfolio choices and access to insurance, and in Section 4, we offer some concluding remarks.

2. The model

In this section, we develop a simple model to highlight how exposure to a large, negative wealth shock affects both household savings and investment. This model shares some key insights with Elmendorf and Kimball (2000).

2.1. Saving and stock investment without insurance

We consider a simple, two-period model in which an investor with an initial wealth of W_0 derives utility from consumption today c_0 and consumption next period c_1. His wealth endowment next period is $\tilde{W}_1$, where, with a small probability $\varepsilon>0$, $\tilde{W}_1=-D<0$, which represents a negative wealth shock, and, with probability $1-\varepsilon$, $\tilde{W}_1=W_1\geq 0$, which represents the present value of all future income.

The investor can consume, save, or invest in a risky asset (stock) at time 0. The initial stock price S_0 is normalized to 1, and the next period stock price is $\tilde{S}_1$, where $\tilde{S}_1=u>1$ with probability $p>0$, and $\tilde{S}_1=d<1$ with probability $1-p$. For simplicity, we assume that the stock price process is independent of the investor's income process and that the interest rate is zero. Then, the investor's problem is to choose consumption c_0, saving α, and stock investment θ to solve

$$\max_{c_0,\alpha,\theta} U(c_0)+\delta E[U(c_1)], \tag{1}$$

subject to the budget constraints

$$\alpha=W_0-c_0-\theta \tag{2}$$

and

$$c_1=\alpha+\theta\tilde{S}_1+\tilde{W}_1, \tag{3}$$

where

$$U(c)=\begin{cases} u(c) & \text{if } c\geq \underline{c},\\ -\infty & \text{otherwise,}\end{cases} \tag{4}$$

where $u(c)$ is strictly increasing and strictly concave for $c\geq\underline{c}$, δ is the subjective time discount rate, and $\underline{c}$ is the subsistence level of consumption. This type of preference is standard in the large literature on habit formation (e.g., Constantinides, 1990; Dybvig, 1995). It essentially assumes that the investor cannot tolerate any consumption below a certain level.

In the above standard setup, we make the following nonstandard assumption.

Assumption A. The magnitude of the negative wealth shock

$$D=W_0-2\underline{c}>0. \tag{5}$$

There are several reasons that such a large negative wealth shock could occur.[5] For example, a sudden occurrence of long-term disability or diagnosis of a terminal condition could imply significant loss of income and large increase in health expenses. Likewise, unanticipated death of a household's primary income earner can represent a significant shock to income, and, in the event of job loss, the required wealth to maintain a minimum standard of living for the uncertain duration of unemployment can often far exceed current wealth (Gruber, 2002).

Using this setup, we derive Proposition 1, regarding optimal portfolio choice.[6]

Proposition 1. *Under Assumption* A, *the investor does not participate in the stock market, i.e.,* $\theta^*=0$. *In addition, the*

[4] While not tested in this paper, a number of empirical studies find that access to health insurance is negatively related to saving rates (e.g., Gruber and Yelowitz, 1999; Chou, Liu, and Hammitt, 2003; Goldman and Maestas, 2005; Qiu, 2006). These findings are also consistent with our hypothesis.

[5] Because D is assumed to be exactly equal to $W_0-2\underline{c}$, it could appear that this assumption is restrictive. However, this assumption is largely for expositional simplicity. First, it is trivial to allow multiple possible negative shock sizes as long as the maximum one is $W_0-2\underline{c}$. Second, if the maximum possible shock exceeds $W_0-2\underline{c}$, then one can replace the negative infinity utility below $\underline{c}$ with a large negative utility and show that the investor still does not participate in the stock market for a wide range of parameter values.

[6] All the proofs are presented in the Appendix.

Author's personal copy

investor consumes only the minimum and saves the rest, i.e., $c_0^* = \underline{c}$ *and* $\alpha^* = W_0 - \underline{c}$.

Proposition 1 says that, no matter how small the (positive) probability of a large wealth shock is or how high the expected stock return is, the investor does not participate in the stock market at all, consumes only at the subsistence level today, and saves the rest to hedge against the possible negative wealth shock in the future. Intuitively, because any loss in the stock investment would incur an infinite utility loss if the bad wealth shock occurs, the investor optimally chooses not to participate in the stock market. Increased consumption or decreased saving poses a similar risk. This result holds for a general class of preferences, wealth profiles, and stock market characteristics. It is also consistent with Carroll (2001), who finds that large shocks can have a significant impact on saving and an even greater impact on risky asset investment, despite the small probability of such shocks.

Nonparticipation and an extremely high saving rate arise no matter how high the investor's wealth is in normal states. As long as the wealth shock in one state of the world is bad enough, even the wealthy people will not participate in the stock market and will save almost all of their wealth for the future. This prediction is consistent with the empirical finding that even wealthy people have a low participation rate (Heaton and Lucas, 1997).

2.2. Role of insurance

We now demonstrate the effect of insurance on investment, consumption and saving decisions. We assume that the investor can purchase insurance against the negative wealth shock and that the insurance premium per dollar coverage is equal to $\lambda \geq 0$. Let I be the dollar amount the investor chooses to insure. Consistent with common insurance practice, we assume that the investor cannot overinsure, i.e.,

$$I \leq D. \tag{6}$$

With the access to insurance, the investor solves

$$\max_{c_0,\alpha,\theta,I} U(c_0) + \delta E[U(c_1)], \tag{7}$$

subject to Eq. (6) and the budget constraints

$$\alpha = w - c_0 - \theta - \lambda I \tag{8}$$

and

$$c_1 = \alpha + \theta \tilde{S}_1 + \tilde{W}_1 + I 1_{\{\tilde{W}_1 = -D\}}. \tag{9}$$

For simplicity, we assume that the insurance premium is fair, i.e., $\lambda = \varepsilon$. Then, it can be shown that full insurance is optimal because the cost of insurance is small relative to the benefit of smoothing consumption across time by consuming more today and investing more in risky stocks. On stock market participation and consumption, we then have Proposition 2.

Proposition 2. *If the stock has nonzero risk premium, then* $\theta^* \neq 0$ *with access to insurance. In addition, given any* $c_0 \in [\underline{c}, W_0 - \underline{c})$, *there exist* $\overline{W}_1$ *and* $\underline{\varepsilon}$ *such that,* $\forall W_1 > \overline{W}_1$ *and* $\varepsilon < \underline{\varepsilon}$, *the optimal time* 0 *consumption* c_0^* *is strictly greater than* c_0.

This proposition suggests that with insurance, as long as the risk premium is nonzero, an investor optimally participates in the stock market, even when he is poor. This result suggests that developing an insurance market could help improve stock market participation and lower risk premium. In practice, insurance coverage and costs vary across products and regions. Proposition 2 implies that the participation rates in the stock market and in the insurance market are positively correlated.

This proposition also implies that access to insurance can significantly increase consumption and reduce saving. For example, if the investor's future wealth is high and the probability of a negative wealth shock is low, then the investor could consume almost all of his initial wealth and save almost nothing when insurance is available. Finally, because participation increases and saving decreases with access to insurance, Proposition 2 predicts a negative correlation between participation and saving rates.

3. Empirical analysis

The model yields a number of testable hypotheses regarding stock market participation rates, saving rates, and insurance coverage. In this section, we test these hypotheses using both cross-country and within-country comparisons. First, we test whether differences in access to private and public insurance across countries and over time are correlated with saving rates and investors' stock holdings. Second, we use US state-level variation in unemployment insurance generosity to test whether unemployment insurance programs are related to observed differences in stock market participation rates across US states. Finally, we test whether US individuals' coverage via an employee-sponsored health, life, or long-term care insurance policy is related to their portfolio choices.

3.1. Country-level analysis

We begin with an analysis of country-level participation rates, insurance penetration, public insurance programs, and saving rates.

3.1.1. Participation and insurance

In our first empirical test, we analyze the correlation between stock market participation levels and the development of private insurance markets across countries. A primary finding of our model is that, in the absence of access to low-cost insurance against large, negative wealth shocks, individuals hold less stocks. Therefore, we should expect to find that in countries with less developed private insurance markets that, all else equal, stock market participation rates should be lower. To test this hypothesis, we run the cross-country regression

$$\text{participation}_c = \beta_0 + \beta_1 \text{insurance}_c + \gamma X_c + \varepsilon_c, \tag{10}$$

where *participation* is the share of individuals in country c that participate in the domestic stock market via direct stock holdings. We obtain our measure of *participation* from Giannetti and Koskinen (forthcoming), who construct country-level participation rates for 26

Author's personal copy

Table 1
Stock market participation and insurance.

This table presents the estimates from the cross-country regression of stock market participation rates onto measures of insurance. The dependent variable, Stock Market Participation Rate, is constructed using country-level surveys conducted between 1997 and 2000 and reported in Giannetti and Koskinen (forthcoming). Private Insurance is measured as the ratio of life and nonlife insurance premiums to gross domestic product as described in Beck, Demirguc-Kunt, and Levine (2000). Social Security Coverage is the average fraction of working-age population covered by a social security program as reported by the World Bank World Savings database, 1987–1993. Other controls are Financial Development, Shareholder Rights, Trust, and Common Law. Financial Development is the ratio of total stock market capitalization to GDP in 1995. Shareholder Rights is measured using the anti-director rights index from La Porta, Lopez-de-Silanes, Shleifer, and Vishny (1998). Trust is the fraction of individuals in country that reply that "most people can be trusted" as reported in the World Values Survey, 1995–1999. Common Law is an indicator equal to 1 if the legal origin of the country is English, as given by La Porta, Lopez-de-Silanes, Shleifer, and Vishny (1998). Bootstrapped standard errors are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

	Stock Market Participation Rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Private Insurance	2.18***		2.21***	1.91***	1.77***	2.09***	1.92**
	(0.42)		(0.52)	(0.48)	(0.38)	(0.44)	(0.81)
Social Security Coverage		0.29**					
		(0.14)					
Financial Development			0.006				−0.045
			(0.047)				(0.063)
Shareholder Rights				0.027***			0.012
				(0.010)			(0.019)
Trust					0.33***		0.29**
					(0.09)		(0.13)
Common Law						0.060*	0.063
						(0.036)	(0.067)
Number of observations	24	14	24	24	23	24	22
R^2	0.39	0.68	0.40	0.53	0.57	0.47	0.71

countries based on surveys conducted between 1997 and 2000.[7] The variable *insurance* measures the degree of private insurance in country *c*. This is measured using the ratio of life and nonlife insurance premiums to gross domestic product (GDP) in 1995 as reported in the Financial Structure database produced by the World Bank and described in Beck, Demirguc-Kunt, and Levine (2000). A larger *insurance* value would indicate a relatively more developed private insurance market in that country. All subsequent results are also robust to measuring insurance penetration in other years or to using instead an average of insurance penetration in previous years. Finally, X_c includes a variety of country-level controls discussed below in more details, and, because of our small sample size, we bootstrap the standard errors to ensure that we do not overstate the statistical significance of our findings (Horowitz, 2001). The estimates from this regression are reported in Table 1.

Consistent with our model's implications, we find a strong positive correlation between stock market participation rates and private insurance penetration. The univariate regression results reported in Column 1 of Table 1 show a strong and statistically significant relation between private insurance markets and stock market participation levels, and the magnitude of the correlation is economically significant. An increase in a country's extent of private insurance by one standard deviation is associated with a 7 percentage point increase in the country's participation rate.

In Fig. 1, we provide a scatter plot of the data to demonstrate that the positive correlation between participation rates and private insurance is robust and not driven by outliers. Moreover, no evidence exists that the findings are driven by a particular type of insurance. Though not reported here, we find that the strong positive correlation is statistically significant for both life and nonlife insurance. A priori it is unclear which measure of insurance better captures an individual's ability to insure against large, adverse wealth shocks. Life insurance plays an important role in insuring households against the death of the primary income earner in the household, while nonlife insurance products such as health, automobile, and home insurance could provide valuable protections against other adverse income shocks. Therefore, in subsequent analysis, we continue to use

[7] See Giannetti and Koskinen (forthcoming) for more details about how the participation rates were calculated and the various dates at which the surveys were conducted. These data were also used by Guiso, Sapienza, and Zingales (2008).

Author's personal copy

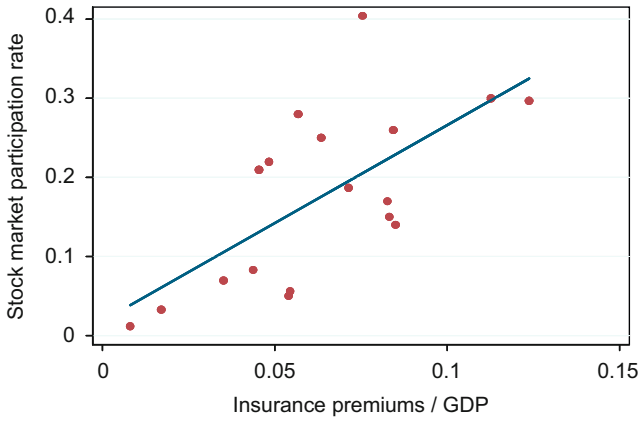

Fig. 1. Stock market participation rates and insurance. This figure displays the univariate scatter plot of stock market participation rates across countries and the ratio of insurance premiums to gross domestic product in 1995. Participation rates are constructed using country-level surveys conducted between 1997 and 2000 as reported in Giannetti and Koskinen (forthcoming). The ratio of insurance premiums to GDP is obtained from the Financial Structure database produced by the World Bank and described in Beck, Demirguc-Kunt, and Levine (2000). The line represents the fitted estimates obtained from the univariate regression.

the combination of life and nonlife insurance premiums over GDP to better capture the overall development of a country's private insurance market.

Our model also suggests that government safety nets, such as unemployment insurance, disability insurance, and publicly funded social security programs, should have an impact on household portfolio choices. By eliminating the possibility of suffering from a large, negative shock to income, these programs could induce individuals to increase their holdings of risky stocks and reduce their overall saving rate.

To test this hypothesis, we replace our measure of private insurance with the percentage of individuals in the country who are covered by a government-funded social security program. By providing a base amount of income for individuals in retirement, a social security program implicitly insures individuals against large adverse shocks to their pre-retirement income. We acquire information about a country's social security program through the World Bank's World Savings database, which reports the fraction of a country's working-age population covered by a social security program. The data set covers 86 countries over the years 1975–1994, though the coverage for most countries is limited to just two or three observations between the years 1987–1993. For the 14 countries that have both stock market participation and social security coverage data available, we calculate a country's average social security coverage rate reported by the World Bank's savings data.

The estimates suggest a very strong positive correlation between social security coverage and stock market participation rates in a country. As shown in Column 2 of Table 1, we find that an increase in a country's social security coverage by 1 percentage point is associated with a stock market participation rate that is 0.29 percentage points larger.

One weakness of the cross-country comparisons, however, is that the correlation between insurance and stock market participation levels could be driven by omitted factors and might not necessarily capture a causal relation. For example, countries with greater economic or financial development could have both higher participation rates and greater private insurance penetration. Moreover, a number of recent papers have demonstrated the potential importance of shareholder rights, legal origin, and trust for stock market participation rates, and, if these factors are also related to the development of insurance markets, there is an omitted variable bias. To address this concern, a variety of controls are added to the base specification for private insurance, and these results are reported in Columns 3–7 of Table 1.

Further supporting the potential role of insurance markets, the strong positive correlation between stock market participation rates and private insurance continues to hold even after controlling for a country's level of financial development. This is shown in Column 3, where we control for financial development using the total stock market capitalization normalized by GDP. The strong positive correlation between participation rates and private insurance is unchanged, and, surprisingly, no evidence exists that financial development is related to stock market participation rates. In an unreported test, we also tried controlling for financial development using the stock market turnover ratio and the value of shares traded normalized by GDP. Again, the correlation between insurance penetration and participation rates remains largely unchanged, and there is little evidence that financial development is related to participation rates. Subsequent findings are also robust to using alternative definitions for financial development.

The strong, positive correlation between stock market participation rates and private insurance continues to hold even after controlling for a number of other country-level factors that have been demonstrated to be important factors in stock market participation levels. In Column 4 of Table 1, we control for the level of investor protection using *shareholder rights*, which is measured using the anti-director rights index constructed by La Porta, Lopez-de-Silanes, Shleifer, and Vishny (1998). Similar to Giannetti and Koskinen (forthcoming), we find that *shareholder rights* is positively related to participation levels, and the correlation between insurance penetration and participation rates declines only by 14% compared with the unconditional correlation reported in Column 1.[8] This suggests that private insurance penetration is positively correlated with participation rates even after controlling for shareholder rights.

In Column 5, we control for the fraction of individuals in a country replying that "most people can be trusted" in

[8] Kimball, Sahm, and Shapiro (2008) show that because controls are often proxies measured with error, it is important to also demonstrate that the addition of controls does not change the coefficients significantly. The magnitude of the unconditional correlation reported in Column 1 of Table 1 would, however, need to decline by approximately 64% for the coefficient to become no longer statistically significant at the 10% significance level.

Author's personal copy

the World Values Survey, 1995–1999. Guiso, Sapienza, and Zingales (2008) argue that individuals' trust of others is positively related to their willingness to hold stocks, and our estimates confirm their findings. However, insurance penetration still remains a strong predictor of country-level participation rates, and the correlation declines only by 19% relative to the unconditional correlation reported in Column 1. Similarly, the inclusion of controls for legal origin (Column 6), as captured by an indicator for common law origin, also does not alter the findings. Finally, simultaneous inclusion of all controls (Column 7) does not affect the findings. *Insurance* remains a strong, positive predictor of stock market participation levels.

The positive correlation between social security coverage and participation rates, as reported in Column 2 of Table 1, also remains largely unchanged with the inclusion of additional controls. In unreported regressions, we also tried controlling for economic development, using the log of GDP per capita in 1995. Again, the correlation between *private insurance* and participation rates remains largely unchanged, while economic development was not significantly related to participation rates. In addition, controlling for the fraction of a country's shares that are "closely held", as measured by Dahlquist, Pinkowitz, Stulz, and Williamson (2003), does not alter the findings.

3.1.2. Insurance and saving rates

While the strong correlation of stock market participation rates with measures of private and public insurance coverage supports our argument that insurance could play a crucial role in explaining participation rates, the findings still could be driven by other country-level determinants not controlled for in these regressions. One potential way to address this concern would be to use variation in the development of insurance markets and stock market participation levels over time along with country-level fixed effects, but, unfortunately, the lack of data on stock market participation over time makes this type of analysis infeasible.

The model, however, does suggest another important implication pertaining to saving rates that yields itself to such a test. Specifically, the model suggests that less-developed countries could better achieve their goal of decreasing their private saving rate and increasing consumption by developing their private insurance industry. Because there exist savings data and insurance penetration data for a number of countries over time, it is possible to conduct a more rigorous test of this prediction that eliminates time-invariant determinants via the inclusion of country-level fixed effects. To test this second hypothesis, we run the country-level, fixed effects specification

$$\text{saving rate}_{c,t} = \beta_0 + \beta_1 \text{insurance}_{c,t} + \alpha_c + \delta_t + \varepsilon_{c,t}, \qquad (11)$$

where *saving rate* is the private saving rate of country c in year t and *insurance* is the ratio of private insurance premiums to GDP. For our measure of savings, we rely on the World Savings database, which measures a country's private saving rate as the difference between gross national saving and public sector saving normalized by gross national disposable income.[9] These data go back as early as 1960 for some countries and end in 1995. For our measure of *insurance*, we again use the World Bank's Financial Structure database, which measures total life and nonlife insurance premiums over GDP beginning in 1987. After combining the two data sets on saving rates and insurance premiums, we are left with a panel of 49 countries that report both variables, and these observations range between the years of 1987 and 1995.

Because this regression includes country-level fixed effects, the parameter of interest, β_1, is estimated only using within-country variation in the saving rate over time and captures how an increase in insurance penetration at the country-level is correlated with changes in the private saving rate within that country. The inclusion of country fixed effects excludes the possibility that the findings are driven by time invariant differences in institutions, legal origin, etc. To control for the possibility of a global trend in savings and insurance penetration during our sample time period, we also include year dummies, δ_t. The estimates, with standard errors clustered at the country level, are reported in Table 2.

The estimates strongly support the hypothesis that greater access to private insurance reduces the saving rate of individuals. Using the entire sample of countries, we find that an increase in private insurance penetration is strongly and negatively correlated with the overall saving rate, as reported in Column 1 of Table 2. The finding holds even after controlling for GDP per capita and overall financial development, as shown in Columns 2 and 3. While GDP per capita is not strongly correlated with saving rates, a strong, negative correlation exists between financial development and saving rates. This is consistent with financial development providing individuals greater ability to transfer resources across time so as to smooth consumption and reduce their need for precautionary savings.

A greater level of public insurance, as captured by a country's fraction of public expenditures on social security (as reported by the World Bank's World Saving database), is also negatively correlated with private savings rates. This is seen in Column 4 of Table 2. The magnitude of the coefficient indicates that a 1 standard deviation increase in the fraction of public expenditures for social security is associated with a decrease of 13 percentage points in the country's private saving rate.

The model also suggests that the negative relation between insurance availability and saving rates should be particularly acute when individuals are more likely to experience income shocks that put them below a subsistence level. Therefore, we might expect to find the negative relation between insurance and saving rates to be particularly strong in less-developed countries where such shocks are likely more common. The evidence appears to support this hypothesis. When we split the sample into less- and more-developed countries, we find

[9] Public sector savings is defined as the savings by the central, state, local, and regional governments plus state-owned enterprises and other nonfinancial public enterprises. See the World Savings database help file for Module 3A, Revision 3.00, on how this measure is constructed.

Author's personal copy

Table 2
Saving rates and insurance.

This table presents the estimates from the country-level, fixed effects regression of private saving rates onto measures of insurance and year dummies. The dependent variable, Private Saving Rate, is constructed using the ratio of private saving to gross national disposable income reported by the World Savings database, where private savings represents the difference between gross national savings and public sector savings. Private Insurance is the ratio of life and nonlife insurance premiums to gross domestic product as reported by the Financial Structure database produced by the World Bank. Social Security Expenses is the fraction of public expenditures spent on social security as reported by the World Bank World Savings database. Other controls are Log(GDP per capita) and Financial Development. GDP per capita is calculated using the International Monetary Fund's International Financial Statistics database, and Financial Development is the ratio of stock market capitalization to GDP reported in the financial structure database. In Column 5, the sample is restricted to countries classified as "low income" and "lower middle income" by the financial structure database, and Column 6 restricts the sample to all other countries. Standard errors are clustered at the country level and reported in parentheses. *** and ** indicate statistical significance at the 1% and 5% level, respectively.

	Private Saving Rate					
	All countries				Low-income	High-income
	(1)	(2)	(3)	(4)	(5)	(6)
Private Insurance	−0.299***	−0.254***	−0.286***		−0.321***	−0.128
	(0.043)	(0.049)	(0.054)		(0.081)	(0.491)
Social Security Expenses				−0.884***		
				(0.080)		
Log(GDP per capita)		−0.037	−0.011	0.037	0.026	−0.029
		(0.032)	(0.026)	(0.027)	(0.051)	(0.029)
Financial Development			−0.027**	0.080	0.041	−0.025**
			(0.011)	(0.063)	(0.049)	(0.011)
Countries	49	49	44	25	11	33
Number of observations	333	333	245	199	65	180
R^2	0.69	0.69	0.81	0.74	0.80	0.83

that the negative correlation between private insurance penetration and saving rates is predominately driven by developing countries. This is seen in Columns 5 and 6, where we divide the sample between "low-income" and "high-income" countries, where countries are classified as "low-income" if the financial structure database lists their income group as either "lower-income" or "lower-middle-income", and as "high-income" otherwise. While the correlation between saving rates and insurance penetration is negative in both samples, it is statistically significant only among the less-developed countries. In unreported results, we also find suggestive evidence that the negative correlation between public insurance, as captured by social security expenses, and private savings rates is larger in less-developed countries. This stronger correlation in low-income countries suggests the negative correlation between social security expenses and savings rates shown earlier is not driven entirely by social security acting as a direct substitute for personal savings by individuals.

3.1.3. Participation and saving rates

Another implication of our theory is that high saving rates should be associated with low participation rates.[10] This negative correlation between saving rates and participation rates runs counter to theories based primarily on fixed participation costs. If fixed costs were the primary driver of low participation rates, a high saving rate would likely generate a larger fraction of people with enough wealth to justify participation, which would instead imply a positive correlation.

To analyze the correlation between participation and savings rates, we now regress country-level participation rates on private savings rates. In particular, we estimate the cross-country regression

$$\text{participation}_c = \beta_0 + \beta_1 \text{saving rate}_{c,t} + \gamma X_c + \varepsilon_c, \quad (12)$$

where *participation* is the share of individuals in country *c* that participate in the domestic stock market via direct stock holdings as defined earlier, and *saving rate* is the average private saving rate of country *c* from 1990 to 1995 calculated using the World Savings database. We again bootstrap the standard errors to reduce the likelihood of overstating the statistical significance of our estimates, which are reported in Table 3.

As suggested by our theory, the univariate correlation between participation rates and saving rates is negative. This is seen in Column 1 of Table 3. The economic significance is large as well. The estimated coefficient indicates that an increase in a country's average private saving rate by 1 standard deviation is associated with a 0.5 standard deviation reduction in a country's participation rate. The evidence is inconsistent with theories based primarily on fixed participation costs.

[10] We thank Miles Kimball for pointing out this result and suggesting the following tests.

Author's personal copy

Table 3
Stock market participation and saving rates.

This table presents the estimates from the cross-country regression of stock market participation rates onto private saving rates. The dependent variable, Stock Market Participation Rate, is constructed using country-level surveys conducted between 1997 and 2000 and reported in Giannetti and Koskinen (forthcoming). Private Saving Rate is constructed using the average ratio of private saving to gross national disposable income reported by the World Savings database from 1990 to 1995, where private savings represents the difference between gross national savings and public sector savings. Other controls include Log(GDP per capita) and Financial Development. GDP per capita is calculated using the International Monetary Fund's International Financial Statistics database, and Financial Development is the ratio of stock market capitalization to gross domestic product reported in the Financial Structure database. Bootstrapped standard errors are reported in parentheses. *** and ** indicate statistical significance at the 1% and 5% level, respectively.

	Stock Market Participation Rate		
	(1)	(2)	(3)
Private Saving Rate	−1.28***	−1.32**	−1.61***
	(0.462)	(0.642)	(0.509)
Log(GDP per capita)		0.049	0.041
		(0.039)	(0.029)
Financial Development			0.115
			(0.098)
Number of observations	19	18	17
R^2	0.26	0.48	0.62

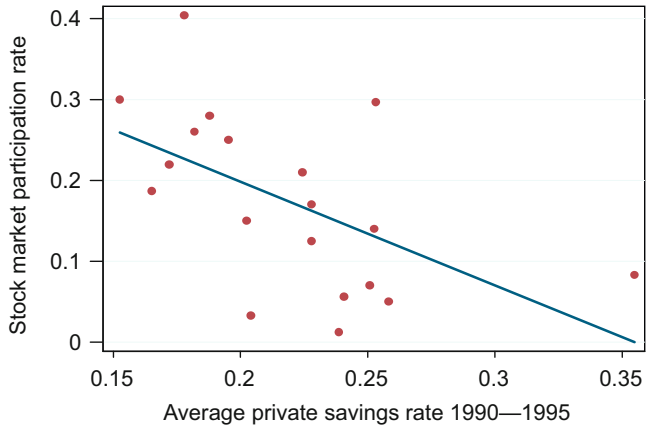

Fig. 2. Participation rates and saving rates. This figure displays the univariate scatter plot of stock market participation rates across countries and the average private saving rates from 1990 to 1995. Participation rates are constructed using country-level surveys conducted between 1997 and 2000 as reported in Giannetti and Koskinen (forthcoming). The private saving rate is constructed using the ratio of private saving to gross national disposable income reported by the World Savings database, where private savings represents the difference between gross national savings and public sector savings. The line represents the fitted estimates obtained from the univariate regression.

The negative correlation between participation and saving rates is also robust and not driven by outliers, as shown in Fig. 2. Including controls for income, as captured by GDP per capita, and financial development do not affect the magnitude of the correlation. This is seen in Columns 2 and 3 of Table 3.

3.2. Participation and unemployment insurance in the US

Overall, the country-level analysis demonstrates that access to insurance correlates positively with stock market participation and negatively with saving rates. However, the availability of insurance could differ across countries in a systematic way that biases our estimates. While it is encouraging that the estimates remain strongly significant even after controlling for numerous country-level factors and after using only within-country variation over time, it would be helpful to demonstrate that the importance of insurance still holds for a more uniform set of economies and at the household level.

To this effect, differences in unemployment insurance premiums across US states provide a potentially interesting test of our model that suffers less from these concerns. As noted by Gruber (1994) and Chetty (2008), the degree to which unemployed individuals' wages are replaced varies significantly across US states. Given this, our model would suggest that, all else equal, in those states with greater unemployment insurance generosity, fewer individuals are subject to income shocks that drop them below their subsistence level, and, hence, average stock market participation levels should be higher. In addition, while there are certainly differences in household characteristics and local governments across US states, it seems reasonable to expect that these differences are less than those of our country-level analysis, and the greater availability of data at both the household and state level within the US allows us to better control for these differences.

To test the hypothesis that greater unemployment insurance generosity should increase stock market participation levels, we run the regression

$$\text{participation}_{s,1992} = \beta_0 + \beta_1 \text{UIG}_{s,1992} + \varepsilon_s, \quad (13)$$

where *participation* is the fraction of households in state *s* that owned stocks in 1992, as determined by the Health and Retirement Study (HRS) in 1992 and reported by Hong, Kubik, and Stein (2004).[11] *UIG* is our measure of the unemployment insurance generosity for each state. This is measured using the log value of average weekly unemployment benefits paid to individuals in state *s* in 1992, as reported by the US Department of Labor.[12] A higher *UIG* indicates that a state's unemployment insurance program provides relatively higher wages to individuals who are unemployed. As before, we use bootstrapped standard errors.

[11] As noted by Hong, Kubik, and Stein (2004), the HRS survey measure of whether households hold stocks does not include holdings in mutual or retirement funds and, therefore, likely understates the true participation rate.

[12] These average weekly benefits were calculated by Chetty (2008) and downloaded from the author's website, http://obs.rc.fas.harvard.edu/chetty/. Our later regression results are also similar when we use unemployment insurance generosity as measured in other years and when we directly use weekly benefits instead of logged values.

Author's personal copy

Table 4
Stock market participation and unemployment insurance in the US.

This table presents the estimates from a state-level regression of stock market participation rates onto unemployment insurance generosity. In Column 1, the dependent variable, Stock Market Participation, is the fraction of households in a state that report owning stocks in the 1992 Health and Retirement Study. These rates were obtained directly from Hong, Kubik, and Stein (2004). In Columns 2–4, the dependent variable, Abnormal Stock Market Participation, is the state-level residual from a household-level regression of a stock market participation indicator onto household characteristics such as age, income, wealth, race, etc., as reported in Hong, Kubik, and Stein (2004). The right-hand side variable, UI Generosity, is the logged value of average weekly unemployment benefits received in a state in 1992 as reported by the US Department of Labor and calculated by Chetty (2008). Other state-level controls, obtained from the 1990 US Census, include Log(Average Household Income) and percent of persons with some college education (% with College Education). Bootstrapped standard errors are reported in parentheses. ** and * indicate statistical significance at the 5 and 10% level, respectively.

	Stock Market Participation	Abnormal Stock Market Participation		
	(1)	(2)	(3)	(4)
UI Generosity	0.174**	0.083*	0.096**	0.094*
	(0.071)	(0.046)	(0.049)	(0.053)
% with College Education			−0.104	−0.101
			(0.163)	(0.180)
Log(Average Household Income)				0.002
				(0.039)
Number of observations	35	35	35	35
R^2	0.16	0.12	0.13	0.13

In support of our model, we find a strong positive correlation between average stock market participation levels and unemployment insurance generosity across US states (Table 4, Column 1). This finding is important for two reasons. First, it again demonstrates that social insurance programs could play an important role in determining stock market participation levels. Second, it shows that the strong positive correlation between participation levels and insurance coverage observed in the country-level data also exists within the US. This within-country evidence mitigates the potential concern of omitted factors, such as legal origin and financial development, that was present in the country-level analysis.

However, there could still be an endogeneity problem if differences in US household characteristics across states is correlated with unemployment insurance generosity. To address this, we rerun the above regression using the *abnormal participation rate* constructed by Hong, Kubik, and Stein (2004). This state-level measure captures the residuals from a household-level regression of a stock market participation indicator onto household characteristics such as age, income, wealth, years of education, race, risk tolerance, and an urban versus rural indicator.[13] In other words, this variable captures the amount of each state's average participation rate that cannot be explained by characteristics of the household. By using this as our dependent variable instead, we are better able to exclude the possibility that omitted factors at the household level could drive the correlation between unemployment insurance generosity and stock market participation rates across US states.

[13] See Hong, Kubik, and Stein (2004) for more details on how this measure was constructed.

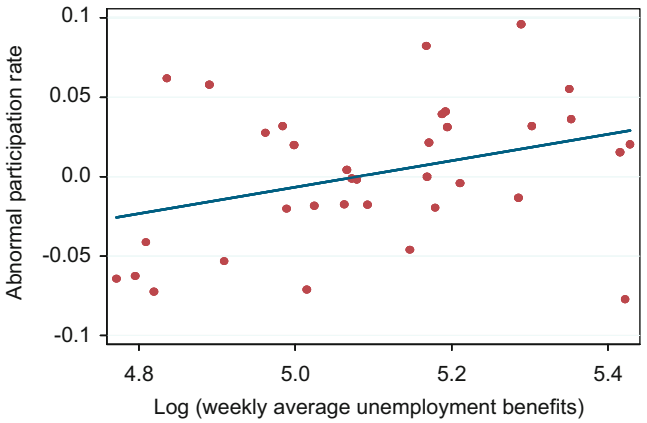

Fig. 3. Participation and unemployment insurance. This figure plots the univariate scatter plot of abnormal stock market participation rates across US states and generosity of unemployment insurance in 1992. Abnormal participation rates are the state-level residuals from a household-level regression of a stock market participation indicator onto household characteristics. These residuals were obtained from Hong, Kubik, and Stein (2004). Unemployment insurance generosity is captured using log(average amount of weekly unemployment benefits paid to individuals in a state), as reported by the US Department of Labor. These weekly benefit averages were obtained from Chetty (2008).

The estimates from this test, as reported in Column 2 of Table 4, indicate the positive correlation between state-level participation rates and unemployment insurance levels is not driven mainly by differences in household characteristics across states. While using the residuals from the household regression weakens the estimate slightly, we still find a positive relation between *UIG* and stock market participation levels that is significant at the 10% level, and the correlation is robust to outliers, as seen in Fig. 3.

Author's personal copy

While the significant correlation between state-level unemployment generosity and participation rates, even after controlling for households characteristics, is highly suggestive of a role for public safety nets, a potential concern still exists that state-level unemployment generosity could be correlated with a state's overall level of economic development or education levels, which could in turn be directly related to participation rates. If this were true, we could be wrongly attributing the impact of overall economic development or education to unemployment generosity.

However, the positive correlation between state-level unemployment generosity and participation rates does not appear driven by the differing levels of development or education across US states. In Column 3, we add a control for the percentage of persons within the state with some college education, and in Column 4, we add a control for the log(average household income) of the state. Both controls are calculated using the 1990 US census. Even after adding these controls, the correlation between participation levels and unemployment insurance generosity remains largely unchanged in magnitude and statistically significant.

Similar to our country-level analysis, we should also expect to find that the importance of insurance for stock market participation should be greater if individuals face greater exposure to large, adverse shocks to their income. In this case, the impact of unemployment insurance should be particularly large in states where a greater fraction of workers are at risk of losing their job. This is, in fact, what we observe in the data. This is shown in Table 5, where we repeat our analysis but include a control for the *risk* of a states' population and the interaction of this term with the unemployment generosity of that state. Our measure of *risk* is an indicator equal to one for states with an above-median fraction of workers in 1992 who are located in two-digit standard industrial classification industries that have seen drops in overall employees in the previous five years. This measure is calculated by using the 1987–1992 annual data reported by the US Census Bureau in its County Business Patterns database. To make the interpretation of our coefficients easier to interpret, we have demeaned *UIG*.

As shown in Column 1 of Table 5, the positive correlation between the generosity of a state's unemployment insurance program and stock market participation levels is driven by states where a large fraction of workers are employed in industries exhibiting significant drops in employment in the past five years. These results are also robust to including controls for percentage of the states' population with at least some college education and the average household income of the state, as seen in Columns 2 and 3. Moreover, it is important to reemphasize that the dependent variable, *abnormal stock market participation*, is the state-level residual from a household-level regression of a stock market participation indicator onto household characteristics. Therefore, the positive impact of unemployment generosity reported in Column 3 is above and beyond what can be explained by household characteristics and state-level measures of education and wealth.

Table 5
Participation, unemployment insurance, and risk.

This table presents the estimates from a state-level regression of stock market participation rates onto unemployment insurance generosity and a measure of how many workers are employed in industries with negative job growth over the past five years. The dependent variable, Abnormal Stock Market Participation, is the state-level residual from a household-level regression of a stock market participation indicator onto household characteristics such as age, income, wealth, and race, as reported in Hong, Kubik, and Stein (2004). The right-hand side variable, UI Generosity, is the demeaned logged value of average weekly unemployment benefits received in a state in 1992 as reported by the US Department of Labor and calculated by Chetty (2008). Risk Exposure is an indicator equal to one for states with an above-median fraction of workers in 1992 who are located in two-digit industries that have seen drops in overall employees in the last five years as calculated using the 1987–1992 annual data reported by the US Census Bureau in its County Business Patterns database. Other state-level controls, obtained from the 1990 US Census, include Log(Average Household Income) and the percent of persons with some college education (% with College Education). Bootstrapped standard errors are reported in parentheses. ** and * indicate significance at the 5 and 10% level, respectively.

	Abnormal Stock Market Participation		
	(1)	(2)	(3)
UI Generosity	0.011	0.013	0.007
	(0.072)	(0.071)	(0.080)
Risk Exposure	0.012	0.003	0.001
	(0.013)	(0.013)	(0.014)
UI Generosity * Risk Exposure	0.157*	0.186*	0.192**
	(0.088)	(0.098)	(0.096)
% with College Education		−0.207	−0.218
		(0.202)	(0.193)
Log(Average Household Income)			0.003
			(0.052)
Number of observations	35	35	35
R^2	0.23	0.25	0.26

3.3. US household participation and health insurance

In addition to potential unemployment, individuals could suffer a large, negative wealth shock following an adverse health shock. Given this, our model also indicates that individuals should be more likely to participate in the stock market when they are insured against such health shocks. This is more likely true when individuals are covered under an employer-sponsored health insurance program. Households could also purchase life or long-term care insurance policies to mitigate such shocks. To that effect, we should expect to find households covered by such insurance policies to be more likely to hold risky stocks.

To test this hypothesis, we estimate the following household regression using Probit:

$$\text{participation}_h = \beta_0 + \beta_1 \text{insurance}_h + \gamma X_h + \varepsilon_s, \quad (14)$$

where *participation* is an indicator that equals one if the household reports owning stocks in the Health and

Author's personal copy

Retirement Study of 1992. *Insurance* includes a number of indicators for whether the household reports being covered by an employer-sponsored health insurance program, life insurance policy, or long-term care policy. Finally, X_h captures a variety of other household characteristics that could affect stock market participation levels such as age, wealth, and income. In particular, we include the variables *Age*, which is the age of the oldest member in the household, and *Education*, which is the largest number of years of education obtained by a household member. We also include logged values of the households' reported wealth (including property) and income as well as indicators *Race* and *Marriage* that equal one if the household reports any nonwhite members or marriage. The marginal effects from this Probit regression are reported in Table 6.[14]

As predicted by the model, we find a strong positive correlation between stock market participation and insurance against adverse health shocks. All three types of insurance (employee-sponsored health insurance, life insurance, and long-term care insurance) exhibit a strong, positive correlation with stock market participation for US households even after controlling for household income, wealth, and education. The magnitudes are economically significant. As reported in Column 4, households covered under an employee-sponsored health insurance program are 10 percentage points more likely to own risky stocks. This finding is consistent with the recent empirical findings of Qiu (2006), that households with health insurance are more likely to own stocks and tend to invest a larger proportion of their financial assets in stocks than uninsured households do, and of Goldman and Maestas (2005), who find that better health insurance is associated with greater investments in risky assets.

These results, however, should be interpreted with caution. While the positive correlation between participation rates and employee-sponsored health insurance is suggestive, a potential endogeneity concern remains. It is possible that households with an employee-sponsored health insurance policy are different in a way which we are not controlling for that makes them more likely to hold risky stocks. There could also be a concern of reverse causality in that households that hold more risky stocks may choose jobs based on the health insurance program offered by employers. While this would still suggest an important relation between insurance and the ownership of risky stocks, the underlying story would be different.

Table 6
Health, life, and long-term care insurance.

This table presents the marginal effect estimates from a US household-level Probit regression of stock market participation rates onto measures of insurance coverage. The dependent variable, Stock Market Participation, is an indicator that equals one if the household reports owning stocks in the 1992 Health and Retirement Study. Employee Health Insurance, Life Insurance, and Long-Term Care Insurance are indicators equal to one if the household reports having a such an insurance policy. Age is the age of the oldest member in the household. Education is the largest number of years of education obtained by a household member. Race is an indicator equal to one if any household member is nonwhite. Marriage is an indicator equal to one if the respondent for the household reports being married. Wealth is the household's total reported wealth including property. Income is the total household income reported. Robust standard errors are reported in parentheses. *** indicates statistical significance at the 1% level.

	Stock Market Participation Indicator			
	(1)	(2)	(3)	(4)
Employee Health Insurance	0.104***			0.102***
	(0.012)			(0.014)
Life Insurance		0.130***		0.105***
		(0.013)		(0.014)
Long-Term Care Insurance			0.127***	0.116***
			(0.047)	(0.048)
Age	0.001	0.001	0.001	0.002
	(0.001)	(0.001)	(0.001)	(0.001)
Education	0.031***	0.031***	0.031***	0.028***
	(0.002)	(0.002)	(0.002)	(0.002)
Race	−0.097***	−0.095***	−0.098***	−0.094***
	(0.013)	(0.013)	(0.013)	(0.013)
Marriage	−0.014	−0.042	−0.019	−0.044
	(0.035)	(0.036)	(0.034)	(0.037)
Log(Wealth)	0.137***	0.133***	0.135***	0.138***
	(0.005)	(0.005)	(0.006)	(0.006)
Log(Income)	0.033***	0.041***	0.051***	0.030***
	(0.009)	(0.009)	(0.009)	(0.009)
Number of observations	9,484	9,427	8,930	8,816
R^2	0.20	0.20	0.20	0.21

4. Concluding remarks

Our empirical analysis finds broad support for the idea that access to insurance is related to households' investment and savings decisions. We find that the development of private insurance markets is positively correlated to stock market participation levels and negatively correlated with savings rates. The degree to which individuals' retirement consumption is insured via a government-funded social security program is also positively related to stock holdings and negatively related to their overall saving rates. These findings are particularly strong in low-income countries where the risk of adverse income shocks is likely more acute. We find similar patterns among US households. Differences in the generosity of unemployment insurance across US states are correlated with households' stock market participation rates, particularly in states where workers are at a greater

[14] The findings in Table 6 are also robust to using an ordinary least squares or Logit specification.

Author's personal copy

risk of losing their jobs. This holds even after controlling for both household characteristics and state-level differences in income and education levels. The use of health, life, and long-term care insurance is also positively related to US household participation rates.

While each of these tests admittedly has their weaknesses, the combination of empirical evidence and the simplicity of our model suggest a potentially important role of insurance for explaining both the limited stock market participation and the low-consumption–high-savings puzzles. To our knowledge, this empirical relation between insurance and household investment and savings decisions has not been shown before.

Our findings suggest that the growth of a private insurance industry or the provision of government-funded social safety nets, such as universal health care, could affect participation rates and domestic consumption. In particular, the dramatic growth of private insurance markets could play a role in explaining reduced savings and increased consumption in many developed countries, including the US. At the same time, the findings indicate that a sudden collapse of the private insurance market (or confidence in contracts that provide insurance to investors) could induce a sharp drop in consumption and a flight to risk-free forms of savings. Such a shift in household portfolios could have implications for the broader economy beyond that of the initial collapse of insurance markets.

The findings also have many implications for developing countries that often exhibit high saving rates and low consumption. Combined, high savings and low consumption increase their dependence on foreign demand for economic growth. To mitigate this dependence, many developing countries adopt various policies attempting to stimulate consumption and lower the domestic saving rate. For example, the Chinese government in recent years lowered its interest rates substantially to dissuade people from saving. Our model suggests that the development of private insurance markets could provide another mechanism through which governments can reduce domestic savings. Moreover, the development of social safety nets could also promote consumption and stock market participation.

Overall, little research is devoted to the impact of insurance on consumption, saving, stock market participation, and asset pricing, and our model suggests that the interactions among insurance, consumption, saving, and portfolio choice could provide a new and promising line of research.

Appendix

In this appendix, we provide all the proofs.

Proof of Proposition 1. The investor's problem can be written as

$$\max_{c_0,\theta} U(c_0)+\delta[p(1-\varepsilon)U(W_0-c_0-\theta+\theta u+W_1) + (1-p)(1-\varepsilon)U(W_0-c_0-\theta+\theta d+W_1) + p\varepsilon U(W_0-c_0-\theta+\theta u-D)+(1-p)\varepsilon U(W_0-c_0-\theta+\theta d-D)]. \quad (15)$$

Suppose $\theta>0$, then we have

$$W_0-c_0-\theta+\theta d-D \le W_0-c_0-\theta+\theta d-(W_0-2\underline{c}) \le \underline{c}-\theta+\theta d<\underline{c}. \quad (16)$$

So the fourth term in Eq. (15) is $-\infty$ and other terms less than $+\infty$. Therefore, we must have $\alpha^*\le 0$. Similar argument on the third term leads to the conclusion that $\alpha^*\ge 0$. Thus we must have $\theta^*=0$. Given that $\alpha^*=0$, and $U(c)=-\infty$ for any $c<\underline{c}$, we must have $c_0^*=\underline{c}$. □

Proof of Proposition 2. Taking the first derivative of the objective function in Eq. (15) with respect to θ and evaluating it at $\varepsilon=0$, we have

$$(pu+(1-p)d-1)[U'(W_0-c_0-\lambda I+W_1)+U'(W_0-c_0-\lambda I+I-D)], \quad (17)$$

which is always positive (negative) if $(pu+(1-p)d-1)>(<)0$ for any optimal choice of $c_0=c_0^*$ and $I=I^*$, because the utility function is strictly increasing. This implies that $\theta^*\neq 0$ as long as the risk premium $(pu+(1-p)d-1)$ is not zero (interest rate is set to zero for notational simplicity).

As W_1 approaches ∞ and ε approaches 0, the first derivative of the objective function in Eq. (15) with respect to c_0 approaches $U'(c_0)>0$. This implies that c_0^* must approach $W_0-\underline{c}$ and hence greater than any $c_0<W_0-\underline{c}$. □

References

Ang, A., Bekaert, G., Liu, J., 2005. Why stocks may disappoint. Journal of Financial Economics 76, 471–508.

Basak, S., Cuoco, D., 1998. An equilibrium model with restricted stock market participation. Review of Financial Studies 11, 309–341.

Beck, T., Demirguc-Kunt, A., Levine, R., 2000. A new database on the structure and development of the financial sector. World Bank Economic Review 14 (3), 597–605.

Campbell, J.Y., 2006. Household finance. Journal of Finance 61, 1553–1604.

Cao, H., Wang, T., Zhang, H.H., 2005. Model uncertainty, limited market participation, and asset prices. Review of Financial Studies 18, 1219–1251.

Cao, S.L., Modigliani, F., 2004. The Chinese saving puzzle and the life-cycle hypothesis. Journal of Economic Literature 42, 145–170.

Carroll, C.D., 2001. Theory of the consumption function, with and without liquidity constraints. Journal of Economic Perspectives 15, 23–46.

Carroll, C.D., Kimball, M., 2008. Precautionary saving and precautionary wealth. In: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics, vol. 6. Palgrave Macmillan, Basingstoke, New York, pp. 579–584.

Carroll, C.D., Overland, J.R., Weil, D.N., 1997. Comparison utility in a growth model. Journal of Economic Growth 2, 339–367.

Carroll, C.D., Overland, J.R., Weil, D.N., 2000. Saving and growth with habit formation. American Economic Review 90, 341–355.

Carroll, C.D., Weil, D.N., 1994. Saving and growth: a reinterpretation. Carnegie-Rochester Conference Series on Public Policy 40, 133–192.

Chamon, M., Prasad, E., 2006. The determinants of household saving in China. Unpublished working paper. International Monetary Fund, Washington, DC.

Chetty, R., 2008. Moral hazard vs. liquidity in unemployment insurance. Journal of Political Economy 116, 173–234.

Chou, S.Y., Liu, J.T., Hammitt, J.K., 2003. National health insurance and precautionary saving: evidence from Taiwan. Journal of Public Economics 87, 1873–1894.

Constantinides, G.M., 1990. Habit formation: a resolution of the equity premium puzzle. Journal of Political Economy 98, 519–543.

Dahlquist, M., Pinkowitz, L., Stulz, R.M., Williamson, R., 2003. Corporate governance and home bias. Journal of Financial and Quantitative Analysis 38, 87–110.

Author's personal copy

Dow, J., Werlang, S., 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. Econometrica 60, 197–204.
Dybvig, P.H., 1995. Duesenberry's racheting of consumption: optimal dynamic allocation and investment given intolerance for any decline in standard of living. Review of Economic Studies 62, 287–313.
Eeckhoudt, L., Kimball, M., 1992. Background risk, prudence, and insurance demand. In: Dionne, G. (Ed.), Contributions to Insurance Economics. Kluwer Academic Press, Boston, pp. 239–254.
Elmendorf, E., Kimball, M., 2000. Taxation of labor income and the demand for risky assets. International Economic Review 41, 801–832.
Epstein, L., Schneider, M., 2007. Learning under ambiguity. Review of Economic Studies 74, 1275–1303.
Giannetti, M., Koskinen, Y., forthcoming. Investor protection, equity returns, and financial globalization. Journal of Financial and Quantitative Analysis.
Goldman, D., Maestas, N., 2005. Medical expenditure risk and household portfolio choice. Unpublished Working Paper 11818. National Bureau of Economic Research, Cambridge, MA.
Gruber, J., 1994. The consumption smoothing benefits of unemployment insurance. American Economic Review 87, 192–203.
Gruber, J., 2002. The wealth of the unemployed. Industrial and Labor Relations Review 55, 79–94.
Gruber, J., Yelowitz, A., 1999. Public health insurance and private savings. Journal of Political Economy 107, 1249–1274.
Guiso, L., Haliassos, M., Jappelli, T., 2003. Household stockholding in Europe: Where do we stand and where do we go? Economic Policy 18, 117–164.
Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. Journal of Finance 63, 2558–2600.
Haliassos, M., Bertaut, C.C., 1995. Why do so few hold stocks? Economic Journal 105, 1110–1129.
Heaton, J., Lucas, D., 1997. Market frictions, savings behavior, and portfolio choice. Macroeconomic Dynamics 1, 76–101.
Heaton, J., Lucas, D., 2000. Portfolio choice and asset prices: the importance of entrepreneurial risk. Journal of Finance 55, 1163–1198.
Hong, H., Kubik, J.D., Stein, J., 2004. Social interaction and stock-market participation. Journal of Finance 59, 137–163.
Horowitz, J.L., 2001. The bootstrap. In: Heckman, J.J., Leamer, E. (Eds.), Handbook of Econometrics, vol. 5. Elsevier Science, New York, pp. 3159–3228.
Huang, J., Wang, J., 2007. Market liquidity, asset prices, and welfare. Unpublished working paper, Massachusetts Institute of Technology, Cambridge, MA.
Kimball, M., 1992. Precautionary motives for holding assets. In: Newman, P., Milgate, M., Eatwell, J. (Eds.), The New Palgrave Dictionary of Money and Finance. Stockton Press, New York, pp. 158–161.
Kimball, M., 1993. Standard risk aversion. Econometrica 60, 589–611.
Kimball, M., Sahm, C., Shapiro, M., 2008. Imputing risk tolerance from survey responses. Journal of the American Statistical Association: Applications 103, 1028–1038.
La Porta, R., Lopez-de-Silanes, F., Shleifer, A., Vishny, R.W., 1998. Law and finance. Journal of Political Economy 106, 1113–1155.
Linnainmaa, J., 2005. Learning and stock market participation. Unpublished working paper, University of Chicago Graduate School of Business, Chicago, IL.
Mankiw, N.G., Zeldes, S.P., 1991. The consumption of stockholders and nonstockholders. Journal of Financial Economics 29, 97–112.
Polkovnichenko, V., 2007. Life cycle portfolio choice with additive habit-formation preferences and uninsurable labor income risk. Review of Financial Studies 20, 83–124.
Qiu, J., 2006. Precautionary saving and health insurance: a portfolio choice perspective. Unpublished working paper, Wilfrid Laurier University, Ontario, Canada.
Vissing-Jorgensen, A., 2002. Towards an explanation of household portfolio choice heterogeneity: nonfinancial income and participation cost structures. Unpublished working paper, National Bureau of Economic Research, Cambridge, MA.
Wall Street Journal Asia, 2009. Geithner's China pitch. June 4.