Center for Quantitative Economics

Investors' favourite – A different look at valuing individual labour income

Jan Voelzke [†], Jeanne Diesteldorf [†], Fabian Goessling [†] und Till Weigt [†]

60/2017

[†] Department of Economics, University of Münster, Germany

wissen.leben
WWU Münster

CQE Working Paper

Investors' favourite -
A different look at valuing individual labour income

Jan Voelzke · Jeanne Diesteldorf ·
Fabian Goessling · Till Weigt

February 13, 2017

Abstract Human capital is a key economic factor in both macro- and microeconomics, and, at least for most people, by far their largest asset. Surprisingly, relatively little effort has been undertaken in the extant literature to empirically determine the value of individual human capital. This paper aims at closing this gap. We use the Substantial-Gain-Loss-Ratio to calculate Good Deal bounds for securitizations of individual labour income one year ahead. Our procedure is applied to US data. We evaluate the attractiveness of hypothetical human capital contracts and can thereby identify investors' favourites.

Keywords Human Capital Contracts, Asset Pricing, Substantial-Gain-Loss-Ratio
JEL: G12, J17, C58

J. Voelzke, Corresponding Author
University of Muenster
Department of Economics
Am Stadtgraben 9
Germany - 48143 Muenster
Telephone: +49 251 8322954
E-mail: mail@janvoelzke.de

J. Diesteldorf · F. Goessling · T. Weigt
University of Muenster
Department of Economics
E-mail: jeanne.diesteldorf@wiwi.uni-muenster.de
E-mail: fabian.goessling@uni-muenster.de
E-mail: till.weigt@uni-muenster.de

1 Introduction

The role of human capital in economics is crucial. However, economic research usually investigates its aggregated form, thereby adopting a purely macroeconomic perspective or rather a viewpoint necessary to examine economy wide asset pricing models. This approach does neither mirror human capital's heterogeneity, nor individual risk sufficiently.

Human capital is, at least for most people, by far the largest asset they possess. In the US, human capital constitutes between 50 and 90 percent of households' overall wealth (Palacios (2015), Baxter and Iermann (1997), Lustig et al. (2013)). If we were dealing with a traditional financial asset, individuals would therefore certainly try to diversify the resulting cluster risk. One (theoretical) possibility to diminish individual income risk is given by human capital contracts, a form of securitizing individual labour income, or other income dependent assets such as certificates on corresponding income indices.

Since income dynamics are risky in the long run, human capital contracts are attractive for individuals, especially for students.[1] The importance of this area for financial investors has been outlined by Shiller (2003) and Huggett and Kaplan (2016). Voelzke (2016) shows that returns on human capital evolve distinctively different to those of stocks for most individuals, making them attractive for investors as a new asset class to diversify their portfolio.[2]

Huggett and Kaplan (2012) offer a thorough review of the relevant literature on the pricing of human capital contracts in general, and on the valuing of aggregated human capital in particular.

In Huggett and Kaplan (2011), the authors derive explicit price bounds for individual human capital by analyzing the joint distribution of financial assets held by individuals and their labour income. They specify an individual stochastic discount factor (SDF) and derive price limits for human capital contracts, by using so called Good Deal bounds.[3] They are based on the assumption that there cannot exist assets which are more attractive than a certain limit. In particular, Hugget and Kaplan restrict the possible Sharpe ratio to narrow price intervals.

As opposed to Huggett and Kaplan (2011), our research adopts the viewpoint of a financial market. Instead of investigating an individual's value for human capital, we aim at determining its hypothetical market value.

[1] See Shiller (2003) for the general advantages for the individual, Palacios (2002) for the special circumstances and historical examples of student financing and Heese and Voelzke (2017) for an example where unexpected technological change in the form of the internet is shown to alter income dynamics significantly.

[2] Further, see Diesteldorf et al. (2016), who emphasize financial investors' need to find new - possibly bubble-free - asset classes to invest in.

[3] Cochrane (2000) develops the corresponding theory and introduces applications.

Methodologically, our approach has two key advantages. Compared to the Sharp-Ratio, employed by Huggett and Kaplan (2011), the Gain-Loss-Ratio which is simply the ratio of expected stochastically discounted gains and expected stochastically discounted losses, has several advantages.[4] More so, our procedure enables the inclusion of a consumption-based model as the underlying pricing mechanism.

We aim at developing an approach that prices assets which behave differently to exchange traded assets. Therefore, other than factor pricing models, such an approach cannot rely on a simple linear combination of existing market prices.

In line with Huggett and Kaplan (2011) and Huggett and Kaplan (2012), our approach captures the key components of individual income dynamics, i.e. an age effect, an individual specific effect, and an idiosyncratically persistent and transitory component. As such, our work reflects the standard of the extant literature on the matter, see Lillard and Weiss (1979) and Guvenen (2009). However, while Huggett and Kaplan (2011) exploit the co-movement of stock returns and aggregated income, our approach captures the occupation specific interdependency via the consumption-based stochastic distcount factors (SDF). We can thereby investigate the heterogeneity between occupational groups which Voelzke (2016) detects for income dynamics of German employees.

In comparison to Huggett and Kaplan (2011), our procedure is more robust, as misspecifications of the underlying models are allowed for and reflected in the price intervals. The occupation specific estimation of the co-movement between SDF and individual labour income explores the differences in attractiveness and increases the pricing and estimation precision. Eventually, we get tighter return intervals, even though we take into account the idiosyncratic risk and the full density.

The two fundamental equations for asset pricing are given by[5]

$$\begin{aligned} p_t &= E(m_{t+1}x_{t+1}) \\ m_{t+1} &= f(\text{data, parameters}), \end{aligned}$$

where p_t equals the asset price at time t, m_{t+1} the SDF, and x_{t+1} the uncertain payoff. Various asset pricing models differ in the definition of their SDF. One can distinguish between models with a fully specified SDF that allow exact pricing for arbitrary assets, and no-arbitrage models, where the SDF is not fully known but is assumed to be consistent with observed prices. In the latter models, prices are only narrowed to no-arbitrage bounds, whereas the former

[4] Bernardo and Ledoit (2000) develop the theory and discuss applications.

[5] Cp. Cochrane (2001) p. xv.

usually fail to empirically fit prices on overall markets.[6]

We strike a balance between the two pricing paradigms by using Good Deal bounds based on the GLR, respectively its advancement, the Substantial-Gain-Loss-Ratio (SGLR).[7] Good Deal bounds are price intervals formed by precluding prices that are too good with respect to some performance measure. In the case of the (S)GLR, the implied prices lie between values given by a fully specified SDF and no-arbitrage bounds. The approach is based on a freely chosen asset pricing model. Attractiveness is then measured depending on the corresponding SDF. We use a SDF of a consumption-based asset pricing model, representing parts of the fundamental pricing mechanism. One of the key advantages of the (S)GLR approach is its ability to incorporate misspecified asset pricing models without losing its validity. The better the SDF is specified, the better the pricing becomes, i.e. the intervals get thinner.

For the actual price calculations we need the joint distribution of both individual income and SDF. We model individual income dynamics and a consumption-based SDF to obtain their joint distributions. Consumption dynamics are described by a VAR model of macroeconomic variables and are included into the income panel model as an exogenous variable. We estimate the model in a Bayesian manner to get the joint predictive posterior, which, in particular, incorporates estimation uncertainty.

An attractiveness limit for asset prices is set by examining the observed financial markets. Eventually, we provide price intervals as Good Deal bounds based on the estimated joint distribution of individual labour income and SDF and the observed SGLR limit. Given the price intervals and the expected payoff, we furthermore calculate expected returns for hypothetical human capital contracts.

Our paper proceeds as follows: Section 2 describes the methodology employed. Section 3 outlines our empirical results, and Section 4 concludes.

[6] See Ludvigson (2011) for an overview of various asset pricing models and their empirical evaluation.

[7] Cp. Bernardo and Ledoit (2000) and Voelzke (2015) concerning both GLR and SGLR.

2 Methodology

We calculate price intervals for a hypothetical human capital contract which securitizes individual labour income for the next year ahead. The resulting uncertain payout is priced by precluding all values that make the generating assets too attractive. Following this idea, we model and estimate the joint behaviour of individual labour income and all factors that influence the chosen measure. Here, this necessitates setting up a model for income and macroeconomic movements. Subsequently, we pick an appropriate attractiveness measure and determine an attractiveness limit on the observed market. Next, we use this limit and the joint distribution of individual labour income and SDF to calculate the Good Deal bounds by the SGLR.

2.1 Determination of an attractiveness limit

One of the recent developments in the research area of Good Deal bounds is the SGLR as developed in Voelzke (2015). It overcomes certain drawbacks of the GLR proposed by Bernardo and Ledoit (2000), which leads to a pricing approach that is regarded as the unification of model based and no-arbitrage asset pricing. In order to calculate the GLR, one needs a suitable SDF, respectively an appropriate asset pricing model. Price bounds based on the GLR can be calculated by finding all prices that imply a GLR smaller than a certain limit. Varying the attractiveness limit from one to infinity corresponds to sliding from pricing based on a fully specified SDF with a unique price to no-arbitrage asset pricing yielding no-arbitrage bounds. We use a consumption-based approach, i.e. the benchmark SDF is based on consumption data.

Following Cochrane (2001), for a simple consumption-based mode, we specify the SDF as

$$m_t \propto (\frac{c_t}{c_{t-1}})^{-3}, \tag{1}$$

where c_t denotes consumption. We use historic market and consumption data to determine the maximally and minimally observed SGLR. Therefore, we calculate the discrete Substantial-Gain-Loss-Ratio (dSGLR) with the algorithm developed in Voelzke and Mentemeier (2016). It is defined as

$$aSGLR^{M}_{\beta,k}(X) \;:=\; \inf_{m' \in aSDF^k_\beta} \frac{\sum_{i=1}^{Tk} (m'_i x_{(\mathrm{i \bmod T})})^+}{\sum_{i=1}^{Tk} (m'_i x_{(\mathrm{i \bmod T})})^-},$$

where $1-\beta$ quantifies the substantial part, k is a grid-thinning parameter, X a vector of payouts, M a vector of corresponding SDFs and $aSDF^k_\beta$ a set of discrete SDFs that are close to M.[8] This number is calculated for different markets before we use minimal and maximal values as an attractiveness limit to price individual labour income.

2.2 The labour income panel model and its estimation

The calculation of price intervals with the SGLR builds on the joint distribution of the SDF and individual labour income. We propose the following model as data generating process

$$\begin{aligned} y_{i,t} &= \alpha_i + \theta age_{i,t} + \delta age_{i,t}^2 + \gamma_k \tilde{c}_t + z_{i,t} + u_{i,t} \\ z_{i,t} &= \rho z_{i,t-1} + \epsilon_{i,t}, \text{ with } \epsilon_{i,t} \sim N(0, \sigma_i^{z2}), u_{i,t} \sim N(0, \sigma_i^{u2}), \end{aligned}$$

where $y_{i,t}$ is the logarithmized labour income of individual i at time t, and $age_{i,t}$ is the age of individual i, and $u_{i,t}, \epsilon_{i,t}$ are error terms that are assumed to be independently identically normally distributed. We use differences in logarithmized consumption $\tilde{c}_t$ to capture its dependency to the SDF.

Furthermore, α_i is an individual parameter determining the general wage level. This parameter picks up all individual properties, such as skills or residence, that affect wage. ρ governs the persistence of the long-run shocks $z_{i,t}$. It is expected to be close to one, since wage processes usually experience several strongly persistent shocks. θ and δ control the influence of age on income. This is common for most models of labour income dynamics.[9] Including the quadratic form captures the observation that wages increase more strongly in the early years of the working life, while typically growing less closer to retirement. Last, γ_k quantifies the co-movement of income and consumption growth, which is key for the proposed pricing procedure. In order to gain statistical power, we model this parameter occupation specific, i.e. $k = 1, \cdots, K$ indexes the occupational affiliation of individual i.

We estimate the model parameters and the density of the joint distribution of $y_{\cdot,t}$, $\tilde{c}_t$, $\tilde{a}_t$ and $\tilde{y}_t$ with a Bayesian approach. For parsimony, we simplify notation by defining the sets of parameters $\alpha = \{\alpha_1, \ldots, \alpha_n\}$, $\sigma^u = \{\sigma_1^u, \ldots, \sigma_n^u\}$ and $\sigma^z = \{\sigma_1^z, \ldots, \sigma_n^z\}$, and $\gamma = \{\gamma_1, \ldots, \gamma_K\}$. Moreover, we define the full parameter set $\theta = \{\theta, \delta, \rho, \gamma, \alpha, \sigma^z, \sigma^u\}$, denote the set θ less θ as $\theta_{-\theta}$, and proceed analogously for all other sets. Thus, the posterior distribution of the parameters given the full set of observations $Y = \{y_{t,i}\}$ with $i = 1 \ldots, n$ and $t = 1, \ldots, T$ is

$$p(\theta|Y) \propto p(Y|\theta)p(\theta). \tag{2}$$

[8] For a detailed explanation and motivation see Voelzke and Mentemeier (2016) and the references therein.

[9] E.g. Lillard and Weiss (1979) and Guvenen (2009).

Its first term is the likelihood of data Y and the second term the joint prior distribution of θ. Note that we assume independent priors, such that the joint distribution can be factorized into one-dimensional distributions.

Our approach to sampling from the posterior distribution (2) does not assume conjugate distributions. Instead, we make extensive use of the Metropolis-Hastings (MH) algorithm. In particular, our sampling approach uses five large Gibbs blocks, such that we iteratively sample from the conditional posteriors

1. $p(\theta, \delta | \theta_{-\{\theta,\delta\}}, Y)$,
2. $p(\rho | \theta_{-\rho}, Y)$,
3. $p(\gamma | \theta_{-\gamma}, Y)$,
4. $p(\alpha | \theta_{-\alpha}, Y)$,
5. $p(\sigma^u, \sigma^z | \theta_{-\{\sigma^u,\sigma^z\}}, Y)$.

Within the last two blocks, we use additional Gibbs steps to sample the individual parameters α_i, σ_i^u and σ_i^z.

To carry out the sampling algorithm, we need to evaluate the likelihood conditioned on different sets of parameters. As the latent variable $z_{t,i}$ does not allow to calculate any likelihood directly, we apply a canonical Kalman filter, which enables us to evaluate any required log-likelihood as the sum of the log predictive densities

$$\log p(Y|\theta) = \sum_{t=1}^{T} \log p(Y_{1:t,\cdot}|\theta, Y_{1:t-1,\cdot}).$$

Note that by the independence of shocks $u_{i,t}$ and $\epsilon_{i,t}$, the joint likelihood of the entire sample is simply the product of individual likelihoods.

As our sampler has to cope with a large number of variables, the proposal distribution of the MH algorithm is of key importance. Therefore, we use an adaptive variant of the MH algorithm, which incrementally scales the variance of a normal proposal such that the acceptance ratio is close to 0.3 for all parameters. We begin to scale the proposal after a burn-in period of $N_b = 10000$ and stop scaling after $N_s = 50000$ iterations of the sampler. We use the subsequent $N = 50000$ draws as our posterior sample. Our estimation results are not driven by the priors as we use the uninformative uniform prior distributions defined in Table 1.

α	θ	γ	ρ	δ	σ^z	σ^u
$\mathcal{U}(-30, 30)$	$\mathcal{U}(0, 10)$	$\mathcal{U}(-20, 20)$	$\mathcal{U}(0.3, 1)$	$\mathcal{U}(-10, 0)$	$\mathcal{U}(0, 5)$	$\mathcal{U}(0, 50)$

Table 1 Prior distributions

Having obtained the posterior parameter draws, we sample from the predictive density

$$p(Y_{T+1,\cdot}|\theta, Y, Z_{T+1,\cdot}),$$

where $Z_{T+1,\cdot} = \{z_{T+1,1}, \ldots, z_{T+1,n}\}$. The latter draws are easily obtained from the Kalman filtering distributions. Whereas the "predicted" age of individuals $i = 1, \ldots, n$ in $T+1$ is straightforward, we additionally require draws from the predictive density of the SDF.

The proposed SDF is based on consumption. To model its (future) distribution as anticipated by investors, we use a simple Baysian VAR approach with Minnesota Prior to obtain a density forecast for consumption.[10] We follow Smets and Wouters (2007) and use US data on GDP, the GDP deflator, the federal funds rate, consumption, investment, hours worked and wages as our variables. The corresponding SDF density can then be sampled $(M_{T+1} := m_{T+1,1}, m_{T+1,2}, m_{T+1,3}, \cdots)$ by using equation (1).

2.3 Deriving Good Deal bounds

Applying the aforementioned procedures results in a sample of the joint distribution of individual labour income and SDF (Y_{T+1}, M_{T+1}), and observed upper and lower attractiveness limits a_u and a_l. Price intervals can now be established by solving the following equations of discrete SGLRs for the lower and the upper price limit p_l and p_u, respectively

$$\begin{aligned} dSGLR_{\beta,k}^{M_{T+1}}(Y_{T+1} - p_l) &= a_l \\ dSGLR_{\beta,k}^{M_{T+1}}(Y_{T+1} - p_u) &= a_u. \end{aligned}$$

Here, β and k are the dSGLR specific parameters described in Voelzke and Mentemeier (2016).[11]

[10] We use a random-walk-in-levels prior for the constant. The freely chosen coefficients of the Minnesota Priors of the parameter covariance matrix are set to 0.5. Cp. Koop and Korobilis (2010) for further details of this approach.

[11] We set $\beta := 0.01$ and $k := 1$, since we prefer to use a large sample from the predictive density instead of using a large k.

3 Empirical results

We use US data to price a hypothetical human capital contract which has a payout in the height of the individual's income in 1998, seen from the viewpoint of and therefore based on information available in 1997.

Interval width and location vary between occupational groups and individuals. Though more robust concerning model misspecification, the price intervals are tighter as in comparable approaches e.g. Huggett and Kaplan (2011).

In the following, we first briefly describe the estimation results for the consumption model and the limit determination, before outlining our outcomes for the income model and the resulting price intervals.

The data for the BVAR model of consumption is taken from Mark W. Watson's Hompage.[12] We use the macroeconomic variables outlined in Smets and Wouters (2007), i.e.: GDP, GDP deflator, federal funds rate, consumption, investment, hours worked and wages in the USA. The federal funds rate is monthly data, whereas the other variables are quarterly data. Except for the federal funds rate, all data is logarithmized. We annualize the data and use observations from 1959 through to 1997 to calculate the density forecast of logarithmized consumption for the year 1998. The median of the predictive density is given by 4.50 and the 2.5%- and 97.25%-quantile are 4.45 and 4.56 respectively, while the true realization in 1998 was 4.51.

Parameter	Median	2.5%−quantile	97.5%−quantile
θ	0.1453	0.1369	0.1533
δ	-0.0010	-0.0011	-0.0009
ρ	0.9770	0.9656	0.9873

Table 2 Summarizing statistics of the posterior of the common parameters.

To obtain yearly individual labour income and occupation, we employ PSID[13] data from 1978 to 1997. We only include individuals with an uninterrupted income trajectory above 6000 Dollar per year that does not exhibit unrealistic outliers.[14] We fix an individual's occupation to be the one that is most often stated over the years. Parameter estimates for the common parameters of the income model can be found in Table 2. Trace plots of the

---

[12] Obtained from `http://www.princeton.edu/~mwatson/ddisk/hendryfestschrift_replicationfiles_April28_2008.zip` in December 2016.

[13] Panel study of income dynamics, public use dataset. Produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI (2016). We use the data in the form stated by the Cross-national Equivalent File (cp. Burkhauser et al. (2000)).

[14] We exclude trajectories that include observations of twice the individual's average income value.

MH procedure and corresponding histograms are shown in Figure 2 in the appendix.[15]

The estimates for θ and δ with positive and negative signs show the expected behaviour. The positive effect of age on income decreases when individuals become older.

The autoregressive parameter ρ is close to one, underlining the long run impact of the persistent shocks. The occupation- and individual-specific parameters vary significantly, mirroring their heterogeneous behaviour.

When calculating the price intervals, we set the maximal attractiveness to 10.5 and the minimal value to 2.6.[16] Both values indicate attractiveness and imply that expected returns will tend to be positive even for assets with a moderate negative correlation to consumption risk.[17] Table 3 summarizes the results for the largest occupational groups of our sample. We calculate return intervals by dividing the expected payout in 1998 by the upper and the lower price bound.

Occ.	median wdth.	lower b.	upper b.	no obs.
archit	11.7%	25.8%	8.6%	32
eng. Tech	14.6%	23.4%	7.8%	13
relatmed	12.3%	19.7%	6.8%	31
mathemat	12.2%	19.5%	7.0%	17
accounta	8.8%	15.6%	5.4%	11
educator	6.7%	16.3%	6.0%	54
scientis	9.1%	15.7%	5.3%	13
security	9.8%	18.8%	7.3%	13
inspecto	12.4%	24.4%	9.3%	18
convey	10.6%	21.7%	7.1%	15
transpor	16.1%	25.1%	9.3%	45
labor/cr	11.2%	19.5%	6.5%	20

Table 3 Occupational group (Occ.), median return interval width (wdth.), average lower (lower b.) and upper return bound (upper b.) and number of observations (no obs.) for occupations with more than ten individuals in the data sample.

[15] Note that the high correlation and autocorrelation of the chains from θ and δ do not harm our analysis. We base our prediction density and SGLR calculation on an i.i.d. draw from the posterior sample, which is equivalent to a thinning factor larger than 100.

[16] This corresponds to the observed values of the 1%-dSGLR for major total return indices of the S&P500. In particular, we investigate the total index returns of the energy, finance, industry, consumer staples, information technology, materials, health care and telecommunication services sector and the S&P500 composite itself between 1989 and 2015. Financial data is taken from Thomson Reuters Datastream; consumption data is provided by the U.S. Bureau of Economic Analysis, retrieved from the hompage of the FRED, Federal Reserve Bank of St. Louis.

[17] E.g. in Figure 5 the estimation result for one author in the data set is given. Even though the posterior of γ is mainly positive, the implied positive correlation with the SDF and the resulting attractiveness is not sufficient to get prices that are larger than the expected payout in 1998. This corresponds to the observation, that on financial markets most assets achieve positive returns on average, even if they posses advantageous dynamics.

Figures 3 - 5 in the appendix visualize estimation results for three exemplary individuals. They differ with respect to their variance value and its decomposition between the persistence and transient shocks, most clearly reflected in their historic income trajectories.

γ differs across occupational groups, reflecting their different exposure to the economic overall movements.

Generally, a negative γ parameter renders a corresponding human capital contract unattractive for a representative investor. By tendency, the corresponding human capital contract pays more in good times and less in bad times. As a result, the corresponding price interval lies more to the left.

The opposite behaviour, i.e. parallel co-movement with the SDF, means that the cash flow is attractive for a representative investor, resulting in intervals that tend to higher prices.

Figure 1 shows a histogram of the logarithmized interval lengths. It visualizes the overall distribution of interval widths and outlines the individual-specific risk on human capital returns.

4 Conclusion

Our paper develops and conducts a new approach to calculating price intervals for individual labour income, all the while accounting for model misspecifica-

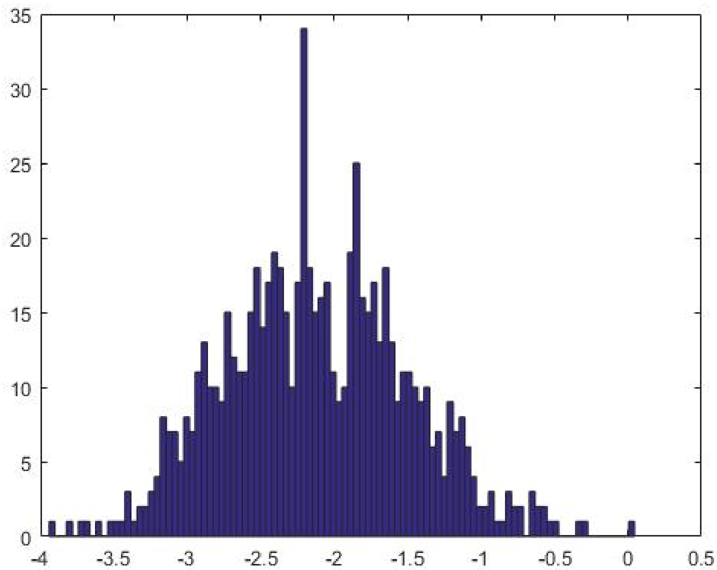

Fig. 1 Histogram of logarithmized return interval width for individuals of the sample.

tion and estimation uncertainty. We incorporate a consumption-based asset pricing approach and adopt the viewpoint of a market representing investor. In particular, pooling by occupational groups enables us to identify the differences in attractiveness of various occupational groups and their individuals as assets for financial investors. Eventually, we state tighter price intervals in comparison to existing approaches in the literature.

Inclusion of an employment dummy into the model to quantify unemployment risk and using a more advanced asset pricing model is left for further research.

References

Baxter, M. and U. J. Iermann (1997). The international diversification puzzle ls worse than you think. *The American Economic Review 87*(1), 170–180.

Bernardo, A. E. and O. Ledoit (2000). Gain, loss, and asset pricing. *Journal of Political Economy 108*(1), 144–172.

Burkhauser, R. V., B. A. Butrica, M. C. Daly, D. R. Lillard, and R. V. Burkhauser (2000). The cross-national equivalent file.

Cochrane, J. H. (2000). Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *The Journal of Political Economy 108*(1), 79–119.

Cochrane, J. H. (2001). Asset pricing, 2001.

Diesteldorf, J., S. Meyer, and J. Voelzke (2016). New evidence for explosive behavior of commodity prices. *Working Paper*.

Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic dynamics 12*(1), 58–79.

Heese, C. and J. Voelzke (2017). Labour income in the age of internet revolution. *The Empirical Economics Letters 16.*

Huggett, M. and G. Kaplan (2011). Human capital values and returns: Bounds implied by earnings and asset returns data. *Journal of Economic Theory 146*(3), 897–919.

Huggett, M. and G. Kaplan (2012). The money value of a man. Technical report, National Bureau of Economic Research.

Huggett, M. and G. Kaplan (2016). How large is the stock component of human capital? *Review of Economic Dynamics 22*, 21 – 51.

Koop, G. and D. Korobilis (2010). *Bayesian multivariate time series methods for empirical macroeconomics.* Now Publishers Inc.

Lillard, L. A. and Y. Weiss (1979). Components of variation in panel earnings data: American scientists 1960-70. *Econometrica: Journal of the Econometric Society*, 437–454.

Ludvigson, S. C. (2011). Advances in consumption-based asset pricing: Empirical tests. Technical report, National Bureau of Economic Research.

Lustig, H., S. Van Nieuwerburgh, and A. Verdelhan (2013). The wealth-consumption ratio. *Review of Asset Pricing Studies 3*(1), 38–94.

Palacios, M. (2002). Human capital contracts.

Palacios, M. (2015). Human capital as an asset class implications from a general equilibrium model. *Review of Financial Studies 28*(4), 978–1023.

Panel study of income dynamics, public use dataset. Produced and distributed by the Institute for Social Research, University of Michigan, Ann Arbor, MI (2016).

Shiller, R. J. (2003). *The new financial order: Risk in the 21st century.* Princeton University Press.

Smets, F. and R. Wouters (2007). Shocks and frictions in us business cycles: A bayesian dsge approach. *The American Economic Review 97*(3), 586–606.

Voelzke, J. (2015). Weakening the gain–loss-ratio measure to make it stronger. *Finance Research Letters 12*, 58–66.

Voelzke, J. (2016). Individual labour income, stock prices and whom it may concern. *Applied Economics Letters*, 1–4.

Voelzke, J. and S. Mentemeier (2016). Computing the substantial-gain-loss-ratio. *Working Paper*.

Acknowledgements

The collection of data used in this study was partly supported by the National Institutes of Health under grant number R01 HD069609 and the National Science Foundation under award number 1157698.

We thank Nicole Branger and Mark Trede for their valuable remarks and suggestions.

Appendix

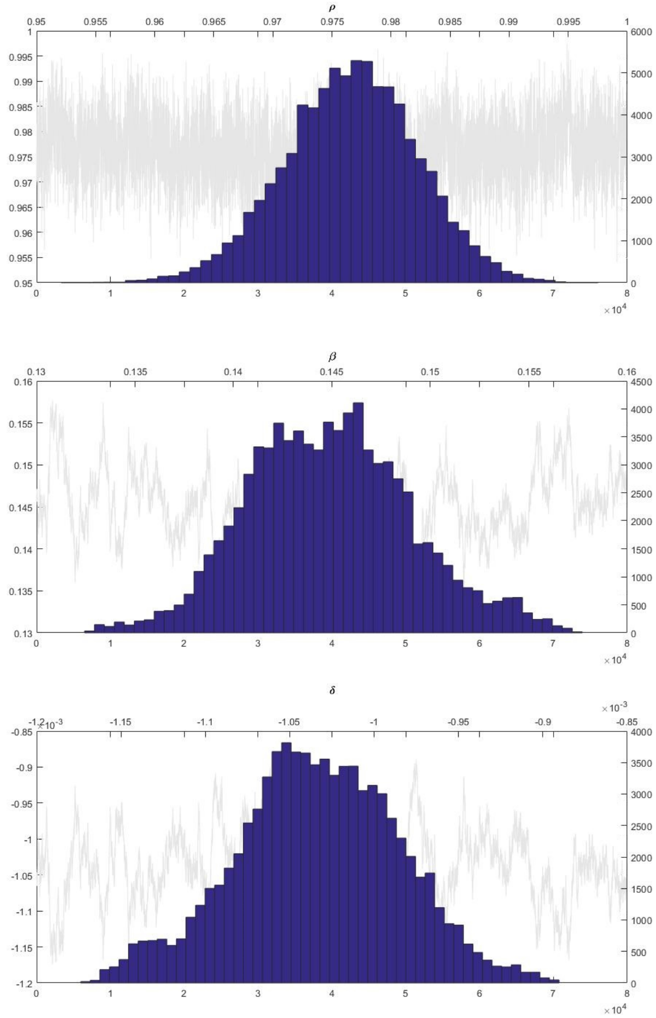

Fig. 2 MCMC results for the common parameters.

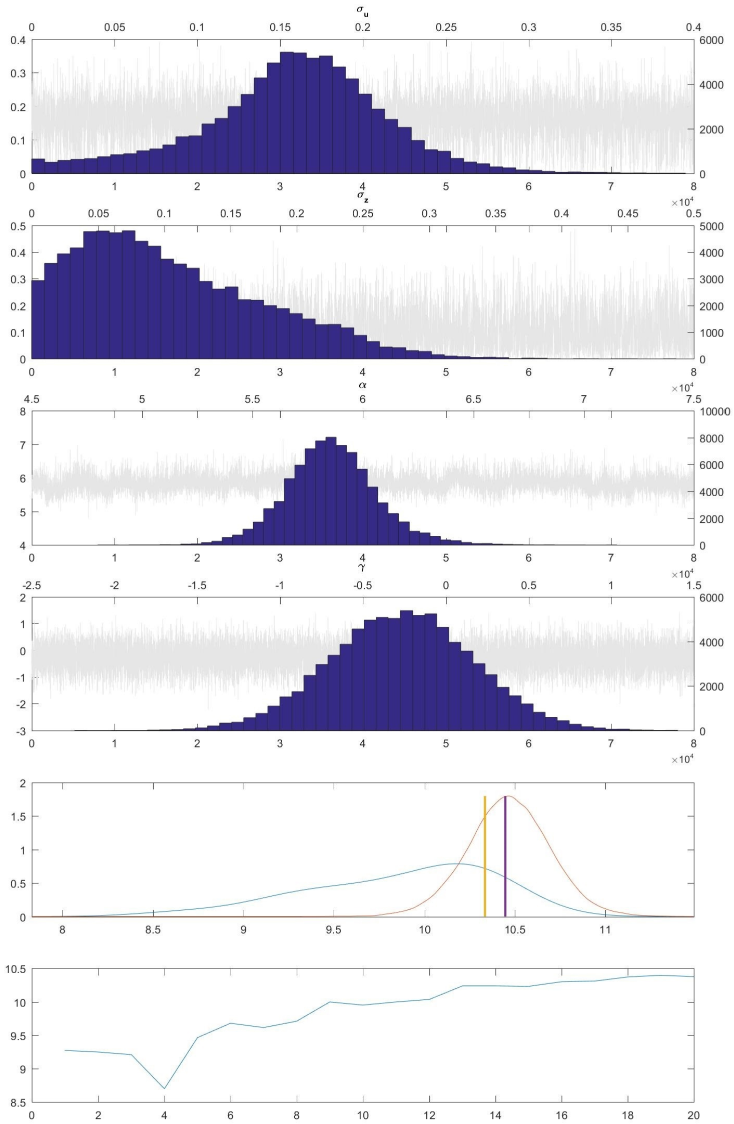

Fig. 3 MCMC results for the parameters of a representative accountant. The first four plots show the MCMC trace plots in gray and a histogram of the posterior for the corresponding parameter. In plot five, the naive density estimate for the historic income trajectory (blue), the density forecast based on the income model (red), and the calculated price intervals are given. At the bottom, the historic income trajectory of the individual is shown.

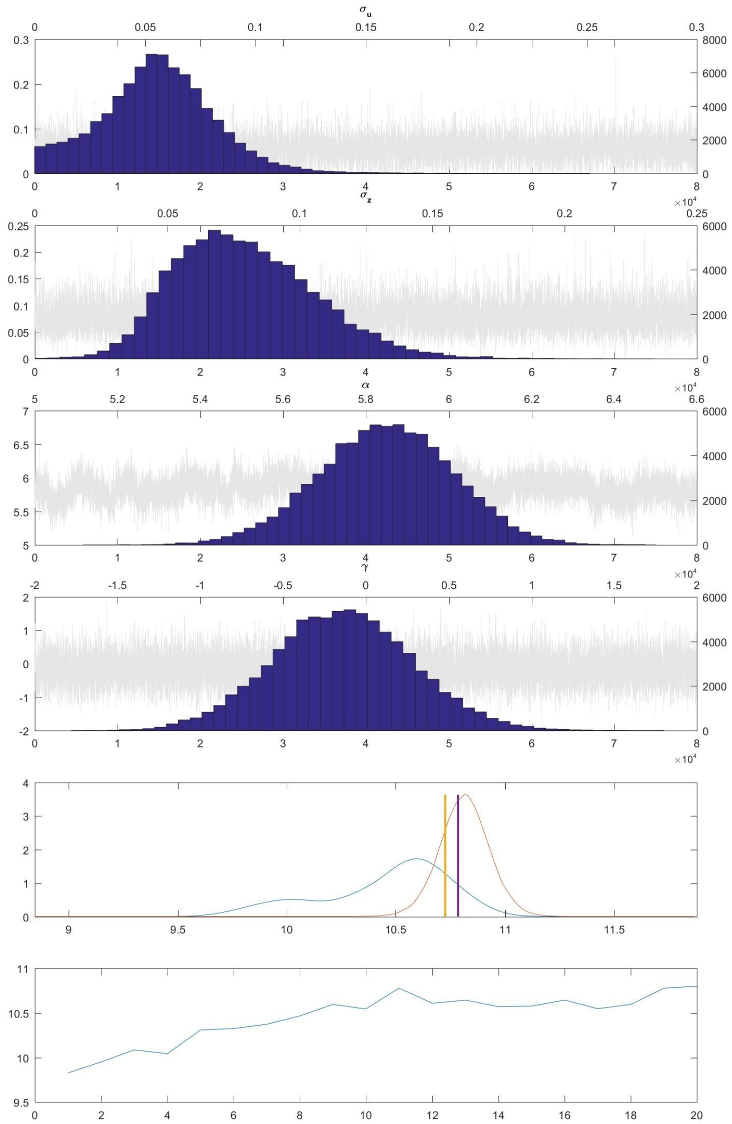

Fig. 4 MCMC results for the parameters of a representative mathematician. The first four plots show the MCMC trace plots in gray and a histogram of the posterior for the corresponding parameter. In plot five, the naive density estimate for the historic income trajectory (blue), the density forecast based on the income model (red), and the calculated price intervals are given. At the bottom, the historic income trajectory of the individual is shown.

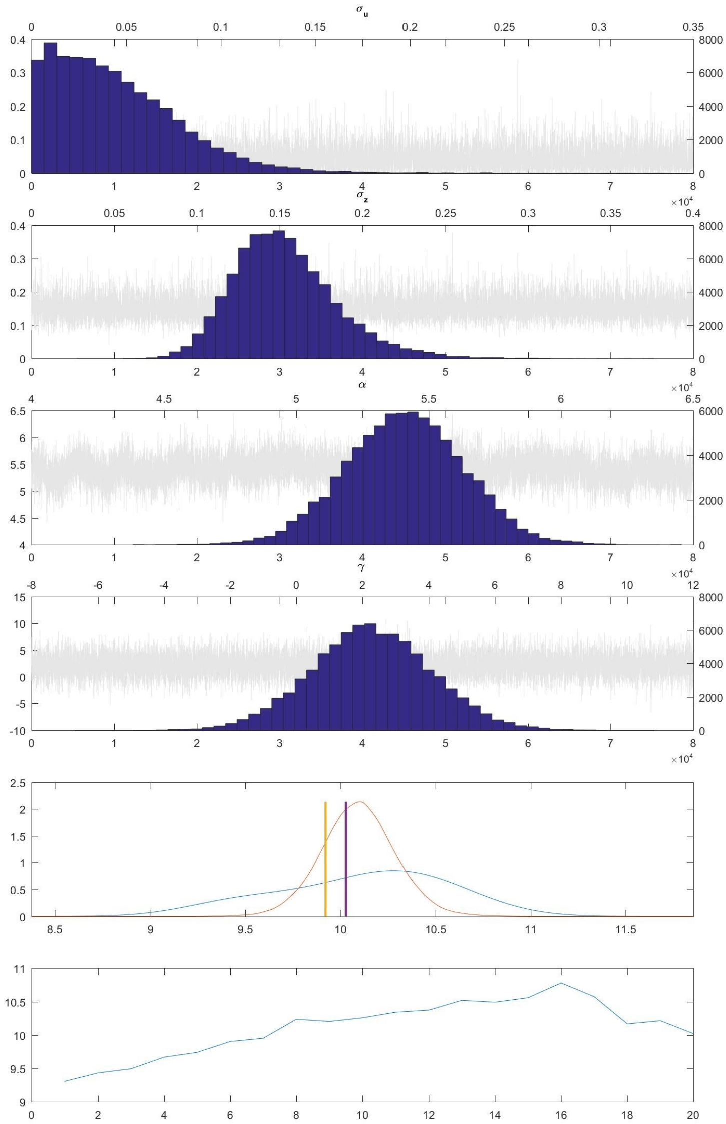

Fig. 5 MCMC results for the parameters of a representative author. The first four plots show the MCMC trace plots in gray and a histogram of the posterior for the corresponding parameter. In plot five, the naive density estimate for the historic income trajectory (blue), the density forecast based on the income model (red), and the calculated price intervals are given. At the bottom, the historic income trajectory of the individual is shown.