

Research statement

Ofer Setty

November 2023

My research focuses on two fields: macroeconomics with a focus on public finance; and finance with a focus on household finance. I use theory and data in both of those fields to study how heterogeneity across various dimensions matters for social policies, and how taking such heterogeneity can help achieving a better design of such policies. In both fields - public finance and household finance – I use models as a lab to experiment with numerous combinations of policy instruments. This allows me to learn about the economic forces at hand, to quantify those forces, and to evaluate the new designs and their potential effects on the economy.

My work is primarily quantitative and structural. This means that the foundations are theoretical while the solution is typically numerical. I use rich models to identify the relevant source or sources of heterogeneity such as age, income, savings and preferences for each unique study. My quantitative work in household finance is typically complemented with an empirical analysis using Scandinavian register data (so far from Sweden and in the near future from Denmark) through my long-time partnership with Scandinavian co-authors. In addition, I use theory whenever possible to highlight the main economic forces at hand. I find that looking at the same question from several methodological perspectives puts more discipline on the analysis and at the same time makes it clearer.

A good example of combining all three methodologies is **Designing Pension Plans According to Consumption-Savings Theory**.¹ The increase in longevity alongside the shift of pension funds from defined benefits to defined contribution systems have increased the importance of the design of pension products. Indeed, financial assets serve as the cornerstone for financial security, and are also the main financial asset for many individuals. In this paper we look at the contribution rates to defined contribution plans. We observe that universally those contribution rates are rigid, as they do not depend on the characteristics of the individual. We challenge this by studying the economic and welfare consequences of relaxing that rigidity.

We start with a theoretical section, where we use a stylized three-period life-cycle model to highlight two guiding principles for optimal savings rates. First, optimal contribution rates depend on expected income growth: the steeper the income profile, the lower the optimal contribution rates for young workers. Plan participants with increasing income over their working life would thus optimally choose contribution rates that increase with age. Second, optimal contribution rates are a decreasing function of the asset balance-to-income ratio.

We then use Swedish registry-based data to test whether consumption-savings behavior is consistent with those principles in the sense that individuals adhere to them in their non-mandated savings outside the pension system.

We find that higher expected income growth is associated with lower savings rates and that average savings rates increase with age. Then, using both OLS and IV approaches to estimate the response of savings rates to shocks to the asset balance-to-income ratio, we find that Individuals reduce their contribution rate after an increase in their asset balance. We then study the reaction to income shocks—the second component of the asset balance-to-income ratio, and find that a

¹ K. Schlafmann, O. Setty and Roine Vestman, “Designing Pension Plans According to Consumption-Savings Theory,” R&R at the Review of Financial Studies.

negative income shock reduces contemporaneous savings. Therefore, the empirical evidence supports the principles for optimal savings rates that we have shown.

Finally, we use a realistic economic environment to design a simple policy rule for contribution rates. We build a quantitative, heterogeneous-agents, life-cycle model that features risky labor income, a pension system with three pillars of retirement savings, and portfolio choice. Our proposal is a contribution rate that depends on the individual's age and balance-to-income ratio. Every year, the contribution rate should unconditionally increase by 0.3 percentage points. Furthermore, investors who fall short by 1 percent from the target balance-to-income ratio for their age should increase their contribution rate by 0.15 percentage points. Our design implies a welfare gain of 1.8 percent and reduces the dispersion of replacement rates by more than 40 percent. These results are robust to time-inconsistent investors.

This work continues the theme of one-size does not fit all that was introduced in my work on pension savings in **On the Design of a Default Pension Fund**.² In 2000 Sweden transformed its traditional retirement payments system into one that includes individual funds for each worker. Using a pro-choice approach Sweden allowed workers to choose among hundreds of investment options, to move at no cost across funds, and to divide their funds across several investment options. Sweden's experience shows that even though workers were encouraged to be active, many chose to stay with the default option, which is a government fund that invests in a well-diversified portfolio. This behavior raises the importance of the design of the default fund, which has consequences for both the risk and the potential gains associated with investment.

We approach the question of the optimal design of the pension fund by developing a model that allows investing the pension's fund in a mix of riskless bonds that have low return and risky stocks that bear high expected return. Our approach focuses on how the policy could take into account the specific characteristics of each investor for customizing the equity share (the share of funds that is invested in stocks, where the rest is invested in bonds). The current frontier with such customizations is to condition the equity share on the age of the investor. A typical rule of thumb for the equity share is 100 minus age: a fund of a 25 year old will have a high equity share of 75%, while that of a 50 years old investor will have a lower equity share of 50%.

We find that two additional characteristics should be taken into account. The first is whether the investor has funds invested in stocks outside the pension account. Those who do not have access to stocks, should have, according to the model, an equity share that is higher by 20 percentage points than those who do hold stocks outside of the pension account. The intuition for this results is that those who already hold stocks outside the pension fund have some of their total assets in stocks, while the others do not. The second additional characteristic is the level of the pension fund. This level allows rebalancing the equity share according to the experience that the investor had with the stock market: an investor who experienced high returns and whose accounts have inflated relative to their income would have lower equity shares than one who had bad experience with the stock market.

Our model indicates that the inclusion of the two additional characteristics via a simple rule of thumb is more important than the existing dependence on the investor's age. Thus, our model not only allows fund managers to better customize the equity share to the specific needs of each

² M. Dahlquist, O. Setty and R. Vestman, "On the Design of a Default Pension Fund," *Journal of Finance*. 2018.

investor in a simple way, but it also opens the door for research for other types of financial customizations that could be useful in this and other contexts.

We supplement our structural study by an empirical one using a dataset that includes a representative sample of the Swedish population. This unique dataset includes, beyond income and socio-economic characteristics, the holdings and the portfolio of investors both in the pension fund and outside the pension fund. We use that data both for studying investors' behavior and as a quantitative input for the calibration of our model.

The balance sheet of households include another important asset, housing, which is the focus of **Macro-prudential Regulations of Mortgage Contracts**.³ As a response to the global financial crisis as well as the prolonged credit boom fueled by low interest rates, macro-prudential policies that directly target households have come into fashion. The purpose of these regulations is to make the economy more resilient to macroeconomic shocks. Similarly, many countries have large mandatory pension systems which force households to save a fixed share of their income each month, regardless of their life situation. The rationale behind this policy is to ensure sufficient income during retirement for everyone and hence to stabilize consumption.

Despite the promised gains in macroeconomic stability, both sets of regulations are controversial because, ultimately, they restrict when and how much households can consume and borrow. Moreover, they tend to put pressure on marginalized groups such as credit constrained households or those with low income. In sum, policy makers must trade off two conflicting goals: macroeconomic stability versus household welfare.

As of now, macro-prudential policies that target households directly is an area where policy makers are ahead of academia in that they pursue policies with little scientific foundation. In this project we aim to develop new macroeconomic models informed by register data from Denmark to explore the mentioned trade-offs and ultimately help policy makers make better informed policy choices.

In this project we use a new quantitative-theoretical macroeconomic heterogeneous agent model, calibrated to register data from Sweden to consider the interaction between various regulation instruments together with their fit for households' characteristics in order to achieve both financial stability and economic welfare.

Heterogeneity across people and in particular their different needs is also present in my work in the field of public finance. In **Universal Basic Income: Inspecting the Mechanisms**, we study Universal Basic Income, a policy that has recently become quite popular among various audiences.⁴ To date, no UBI program has ever been implemented on the aggregate level with a long-term commitment. This lack highlights the need for a quantitative macroeconomic analysis. Yet, the literature evaluating UBI's macroeconomic impact is still at its infancy. In this paper we are interested in analyzing the long-term allocation impact of UBI after the economy has converged to its long-run steady state. We develop a modeling environment that enables us to study a wide range of UBI programs and financing schemes.

³ M. Kilström, K. Schlafmann, O. Setty and Roine Vestman. "Macroprudential Regulations of Mortgage Contracts". Work in progress.

⁴ N. Jaimovich, I. Saporta-Eksten, O. Setty and Yaniv Yedid-Levi, "Universal Basic Income: Inspecting the Mechanisms". 2022. R&R at the Review of Economics and Statistics.

We use a quantitative, production-based general equilibrium model. It is characterized by incomplete markets, individual productivity shocks, and endogenous unemployment and labor force participation. In the model, individuals, make savings and labor market participation decisions. They also face a labor-matching friction as in a standard search-and-matching model. On the government side, we model in detail existing public insurance programs funded by labor and capital distortionary taxation.

We find that the introduction of UBI leads to a large decline in various macroeconomic variables such as output, aggregate capital, and labor force participation. Though UBI programs can reduce inequality and increase consumption for various segments of the population, we find that they have a negative effect on welfare.

UBI depresses economic activity through three main channels. First, financing UBI by increasing distortionary taxation induces a substitution effect, pushing workers out of the labor force. Due to the capital-labor production complementarity in our model, aggregate capital falls as well. We show that the distortionary taxation channel accounts for most of the overall decline in output due to UBI. The remaining impact is split between an "insurance" and an "income effect" channel.

Given the importance of increased distortionary taxation in explaining the drop in labor force participation, we then consider the role of different financing schemes in mitigating UBI's cost. Overall, for each level of UBI, we find that a more progressive income tax scheme mitigates the output costs, and does so by increasing the incentives to join the labor force.

Finally, motivated again by the importance of the labor force participation channel, we consider programs that partially substitute social insurance programs directed at those outside the labor force with UBI. Indeed, a moderate amount of UBI can boost economic activity and welfare, via two channels. First, the direct effect of replacing programs that condition on not participating in the labor force with UBI, which is unconditional, incentivizes labor force participation. Second, given the partial substitution of existing programs in favor of UBI, the tax increase required to finance UBI is smaller.

In some research projects the study of leads to the creation of new policies. Such is the case in **Unemployment Insurance and Unemployment Accounts: The Best of Both Worlds.**⁵ Unemployment insurance is a key policy for insuring workers against unemployment shocks. Its instruments typically include the duration of benefits that are provided to unemployed workers, and the fraction of those payments relative to past earnings (called "replacement rate"). A vast literature characterizes the optimal design of unemployment insurance in various economic environments. This literature aims at striking the right balance between insurance and incentives: on one hand providing benefits insures workers against unemployment and assists workers with smoothing consumption; on the other hand providing benefits may damage the incentives of workers to become employed.

There are only few alternatives to unemployment insurance. One of those, called "unemployment accounts" was implemented in Chile in 2002. This policy is based on individual mandatory savings accounts that accumulate during employment and can be used either during unemployment or at retirement.

⁵ O. Setty, "Unemployment Insurance and Unemployment Accounts: The Best of Both Worlds," *Journal of the European Economic Association*. 2017.

In this paper I move away from fine-tuning unemployment insurance. Instead I design a different type of scheme between the government and the unemployed, which combines features from both unemployment insurance and unemployment accounts. The scheme works as follows. Each worker has an individual account to which she makes monthly deposits during employment and can withdraw from the account at some rate, just like in unemployment accounts. However, once the account is exhausted the worker receives payments from the government as in unemployment insurance. This seemingly simple combination of the two policies have considerable implications for the economy. The main advantage of the proposed scheme is that it allows to *selectively* provide benefits to those who need them the most. In this way the overall tax burden required to finance the system decreases dramatically as many workers insure themselves, while at the same time the government benefits that are provided selectively can be more generous than in unemployment insurance. Thus, the hybrid policy is based on heterogeneity in characteristics of workers that manifest themselves into the level of the individual savings account. These accounts are then used by the government as a proxy for the need of insurance.

In **On the Provision of Unemployment Insurance when Workers are ex-ante Heterogeneous**, we find that unemployment insurance also has an (unintended role) of redistribution across skill groups.⁶ The motivation for this paper is the empirical observation that highly educated workers have lower unemployment rates *and* higher wages. For instance, the median wage of a college graduate is 2.5 times higher than that of a high school dropout, while the unemployment rates for the two groups are around 3% and 9%, respectively. We match these facts in the context of a general equilibrium model with a frictional labor market and incomplete asset markets. Skill heterogeneity is captured by assuming that workers belong to (education) groups that differ in their productivity *and* separation rates. The first is mainly responsible to the differences in wages, while the second matches the differences in unemployment rates.

We use the model to compare the implications with regards to the choice of a replacement rate in models with and without skill heterogeneity. The main quantitative finding is that a model with ex-ante heterogeneity among workers calls for a replacement rate of 33%, more than double the one in a model with ex-ante homogenous workers.

We claim that the reason for this large effect of heterogeneity is the ability of the unemployment insurance system to redistribute resources among groups. To support this claim, we show that the replacement rate falls when either one of the heterogeneity dimensions is (counterfactually) shut down. If there are no productivity differences among workers, then there is only a small gain from redistribution because workers self-insure via savings. When there is no difference in unemployment risk, the UI system has no ability to redistribute resources across groups. Therefore, it is only when the two dimensions of heterogeneity are present that a policymaker has a desire and ability to redistribute using the UI system. To further support this claim we show that the optimal replacement rate further increases when the asset distribution is more dispersed such that high-skill workers also have substantially more wealth, providing more motivation for redistribution. Moreover, the existence of a progressive tax system (calibrated to match the US economy) only marginally lowers the optimal replacement rate. This is because there is still a sufficiently large wealth and consumption dispersion in the economy.

⁶ O. Setty and Y. Yedid-Levi, "On the Provision of Unemployment Insurance when Workers are ex-ante Heterogeneous," The Journal of the European Economic Association, 2020.

My older work touched upon a combination of quantitative analysis together with a deeper theoretical framework, known as *Optimal Unemployment Insurance*. Still, looking at the interplay between workers' heterogeneity and policy design, an importance policy question regards to the assignment of (unemployed) workers to several alternative policies. This link is at the core of my work on Welfare-to-Work programs in **Optimal Design of 'Soft' Welfare Programs**.⁷ Modern economies are characterized by excellent economic conditions and a constant increase in the quality of life. But those economies are also associated with a significant fraction of poor workers that are weakly attached to the labor force. In some countries, such as the United States, those people have been "on welfare", meaning that for years they received welfare payments while being weakly attached to the labor force.

In the last few decades there is a shift from such *welfare* programs to *work* programs where participants are expected to make a considerable effort to become more strongly linked to the labor force. Importantly, those "Welfare-to-Work" programs cater participants with various characteristics that could be potentially important for the success of the program. At the same time there are various activities that participants could be assigned to. Some activities put the emphasis on job search, assuming that once a match is made, the worker will gain experience and talent on the job. Other activities put the emphasis instead on work requirements.

The question then, is how to assign participants to activities in such programs, or in other words, what works best for whom, and why. In this paper we develop a model, where the government can choose among five different policies: *unemployment insurance* that provides workers with benefits while they search on their own; *job-search assistance* that allows a caseworker to assist the worker with finding a job; *mandatory work* that requires the participant to work in a public job that typically creates little value; *transitional work* that combines mandatory work with job-search assistance; and *social assistance* that requires neither search nor work activity from the participant.

We use the optimal contracts approach where we allow the policy to depend on the worker's human capital, her history in the job market and the generosity of the welfare system. We use that framework to learn about the economic forces that shape the optimal design of each of the five policies and the optimal assignment of participants. We find two typical sequences of activities. The first is suitable for participants in economies where the generosity of the welfare system is high. In such economies, a participant that starts with a high human capital is first assigned to job-search assistance. If a job is not found within a given duration, she is assigned to social assistance. Participants in economies with less generous welfare systems start with unemployment insurance and then move to transitional work. At the end of that sequence of activities, the worker is assigned to mandatory work without searching for a regular job.

The intuition for the results works as follows. First, generosity of the welfare system affects the desirability of work activities. The higher the generosity the more expensive it is to compensate the worker for that activity up to a point where work requirements do not pay off anymore. Second, time affects the human capital of the worker. A worker is first given the opportunity to find a regular job. However, as time goes by and her labor market qualifications deteriorate there is little hope for finding a regular job and the worker is relieved from those types of activities. This research sheds light on the efficient assignment of participants to programs. The increased efficiency implies a better use of the economic resources that are designated to those programs.

⁷ N. Pavoni, O. Setty, G. Violante, "Optimal Design of 'Soft' Welfare Programs," *The Review of Economic Dynamics*, 2016.

Perhaps more importantly, that type of research may eventually show the potential of such programs in promoting the integration of the programs' participants to the labor force and to a certain degree to society.

One aspect of unemployment insurance that received little attention in the literature is monitoring the job-search of the worker. This is the focus of **Optimal Unemployment Insurance with Monitoring**.⁸ Here, I analyze a policy that allows the government to collect information about the worker's job-search activities in order to improve the efficiency of the payment scheme for the unemployed worker.

In the classic model of optimal unemployment insurance the job-search effort of the worker is unobserved to the government.⁹ Therefore, the government, who is interested in insuring the worker against unemployment, cannot fully insure the worker as this would damage the incentives of the worker to exert job-search effort. This is where job-search monitoring becomes useful as it allows the government to learn at a cost about the job-search effort of the worker by acquiring a costly signal that is correlated with the worker's job-search effort. The idea is inspired by an actual policy that is practiced in the United States as well as in many other countries where workers are required to provide evidence of their job-search activities. This evidence, such as a list of employment opportunities that the worker has engaged with, can be used by a caseworker at the employment center to validate the effort of the worker, and in some cases where the report is unsatisfactory to sanction the worker in the form of denying a portion of her payments for some period.

The application of monitoring includes two instruments that could be used to improve the efficiency of providing benefits. The first is the accuracy of the signal. More precise signals allow the government to increase the efficiency of the payments scheme. The second is the sanction that is inflicted on the worker if the signal indicates that she did not search for a job. More severe sanctions have a stronger effect for encouraging the worker to search for a job. However, each of those instruments implies a cost for the government. First, the signal's precision is costly. Second, sanctions require the government to compensate the worker for unfair sanctions that result from the noisy signal on the job-search effort.

The existence of such a trade-off between costs and benefits of each instrument motivates the fine balance that maximizes the contract's efficiency. However, in the spirit of my research agenda I do not stop at the *typical* balance between the two instruments. Instead, I continue by showing analytically in a simple model and quantitatively in a rich model how workers' and economies' characteristics affect that balance. The main result is related to a characteristic of workers' preferences (the curvature of the derivative of the inverse of utility) that determines how the cost of compensating workers for sanctions changes with the generosity of the welfare system. Typically, that cost increases with generosity, implying that higher generosity of the welfare system should be associated with acquiring more precise information, and with smaller sanctions.

Occasionally, research opportunities that are unrelated to social insurance policies present themselves. For example, in **Financial Risk and Unemployment** we study how and why aggregate

⁸ O. Setty, "Optimal Unemployment Insurance with Monitoring," *Quantitative Economics*, 2018.

⁹ The main reference here is H. A. Hopenhayn, J. P. Nicolini, "Optimal Unemployment Insurance," *Journal of Political Economy*, 1997. Vol. 105(2), 412–438.

unemployment fluctuates with economic conditions.¹⁰ Unlike productivity and consumption that change only moderately with economic conditions, the unemployment rate fluctuates strongly. For example, during the Great Recession in the United States the unemployment rate peaked at 10% (October 2009), while currently it is little more than 4%. Other recessions have been accompanied by somewhat smaller fluctuations but still very strong ones.

There is a vast literature that tries to explain what drives such strong fluctuations. With a few exceptions the usual suspect in those analyses is productivity. The idea is that lower productivity also lowers the value of employment for both workers and firms, leading to fewer jobs and higher unemployment. This literature, however, has a hard time in explaining how relatively small changes in productivity lead to such strong trends in unemployment.

In this project we take a step back and study interest rates as a mechanism that can explain volatility of unemployment. We document a strong correlation between the interest rate on loans that firms face and the unemployment rate. That correlation is of similar magnitude to that between productivity and unemployment. Motivated by this finding we study the quantitative importance of the interest rate mechanism in a search-and-matching model that is similar to that used for studying productivity shocks, where unemployment responds endogenously to changes in the interest rate that firms face.

We find that the interest rate has a strong quantitative explanation power. We show both analytically and numerically using simulations that the competence of this mechanism comes from the fact that unlike productivity, interest rates are as volatile as unemployment. This research is especially relevant in the wake of financial crises such as the great recession in the United States in 2008-2012.

Another project that regards to pensions savings is **The labor market impacts of mandated savings: Evidence from a Pension Reform in Israel.**¹¹ Here, we study the impact of mandated pension benefits on the labor market. For this purpose we exploit a major pension reform that was gradually implemented in Israel since 2008. The reform required workers and firms to contribute a fraction of the employee's salary to a dedicated individual pension account. We exploit this significant change in the level of pension contributions by examining its impact on various labor market outcomes at both the firm and the worker level.

Our research agenda combines a theoretical model to analyze the effect of mandated pension benefits incorporating heterogeneity of workers and firms, with an empirical study of the equilibrium effect of mandated pension benefits on wages and employment. A key question in this research is the extent to which heterogeneity impacts of mandated benefits according to workers' and firm' characteristics and contrast them with those implied by the model.

This study differs from previous work in the literature of mandated benefits in several ways. First, we will be able to provide a monetary estimate on workers' valuation of pension benefits. A second distinctive feature of our research is that we will provide a detailed examination of heterogeneous impacts according to workers' and firms' characteristics. Our empirical approach will exploit considerable variation in both the extensive and the intensive margins of pension provision which are crucial for the identification of the parameters of interest. The interplay

¹⁰ Z. Eckstein, O. Setty and D. Weiss, "Financial Risk and Unemployment," *International Economic Review*, 2018.

¹¹ A Schlosser, O. Setty and I. Shurtz, "The labor market impacts of mandated savings: Evidence from a Pension Reform in Israel." Work in progress.

between the empirical estimation and the structural modeling will be reflected in two dimensions. First, the theoretical model will be used to gain insights on the economic forces at play that will be tested empirically. Second, the empirical results will then be translated to key parameters that will be fed into the structural model to conduct a welfare analysis.

I conclude this statement with a brief description of **Optimal Social Insurance**.¹² This is an ambitious research agenda that aims at combining various policies in one unified framework in order to learn about the interactions between different policies. Within the incomplete markets literature there is a large literature on social insurance programs, such as unemployment insurance, severance payments, social assistance, universal basic income and taxation. Most of that literature focuses on the optimal design of a given policy, say unemployment insurance, in isolation, sometimes taking other policies into account. Recently, there is a developing literature that studies the optimal design of two policies. Those papers take existing policy instruments and let them compete. Such policy instruments are limited in two important ways. First, within the context of a model they condition on a subset of observables. For example, the design of unemployment insurance is conditioned only on the employment state, taxation conditions only on earnings, and social assistance conditions only on assets. Second, those policies have a rigid structure. For example, unemployment insurance is modeled as a fixed replacement rate for a given period, taxation is modeled parametrically, and social assistance uses a binary transfer (conditional on the assets level). Those two limitations imply that we never study the second best allocation, the one that maximizes welfare given the technology and information structure in the economy. I therefore view this approach as “bottom-up” in the sense that it uses available policies as potential building blocks for maximizing welfare.

In this project I aim at studying the optimal transfer scheme, which I refer to as *Optimal Social Insurance*. This policy is a non-parametric transfer function of all the observables in the model. It therefore overcomes the two limitations described above, and thus, by construction, it achieves the constrained allocation. Analyzing this (somewhat complicated) policy will shed light on the important considerations to be taken into account when designing social insurance. In particular, this policy balances between incentives and insurance in more than one way. On the insurance side I may consider redistribution across types, insurance against unemployment and insurance against wage shocks. On the incentives side I may consider extensive and intensive labor supply, labor demand, and savings.

Once I finish characterizing the optimal social insurance policy I turn to implementation. We consider the economic and welfare effects of each social insurance policy in isolation. Doing so allows me to (1) assess how much of the potential gain achieved by the non-parametric policy can be achieved by each instrument, and (2) understand which dimensions of the insurance-incentives are affected by each instrument. I can then consider several (or all) combinations of instruments. This allows me to understand the important interactions (complementarities/substitution) between the various instruments and to come up with an implementation that gets close to optimal social insurance.

¹² O. Setty, “Optimal Social Insurance.” Work in progress.