



Working Papers of the Priority Programme 1859
Experience and Expectation.
Historical Foundations of Economic Behaviour
Edited by Alexander Nützenadel and Jochen Streb



No 37 (2023, January)

Malmendier, Ulrike

Experience Effects in Economics
Lessons from Past and Current Crises

Arbeitspapiere des Schwerpunktprogramms 1859 der Deutschen Forschungsgemeinschaft
„Erfahrung und Erwartung. Historische Grundlagen ökonomischen Handelns“ /
Working Papers of the German Research Foundation's Priority Programme 1859
“Experience and Expectation. Historical Foundations of Economic Behaviour”

HUMBOLDT-UNIVERSITÄT ZU BERLIN



Published in co-operation with the documentation and publication service of the Humboldt University, Berlin (<https://edoc.hu-berlin.de>).

ISSN: 2510-053X

Redaktion: Alexander Nützenadel, Jochen Streb, Laetitia Lenel

V.i.S.d.P.: Alexander Nützenadel, Jochen Streb

SPP 1859 "Erfahrung und Erwartung. Historische Grundlagen ökonomischen Handelns" Sitz der Geschäftsführung:

Humboldt-Universität

Friedrichstr. 191-193, 10117 Berlin

Tel: 0049-30-2093-70615, Fax: 0049-30-2093-70644

Web: <https://www.experience-expectation.de>

Koordinatoren: Alexander Nützenadel, Jochen Streb

Assistentin der Koordinatoren: Laetitia Lenel

Recommended citation:

Malmendier, Ulrike (2023): *Experience Effects in Economics. Lessons from Past and Current Crises*. Working Papers of the Priority Programme 1859 "Experience and Expectation. Historical Foundations of Economic Behaviour" No 37 (January), Berlin

© 2022 DFG-Schwerpunktprogramm 1859 „Erfahrung und Erwartung. Historische Grundlagen ökonomischen Handelns“

The opinions and conclusions set forth in the Working Papers of the Priority Programme 1859 *Experience and Expectation. Historical Foundations of Economic Behaviour* are those of the authors. Reprints and any other use for publication that goes beyond the usual quotations and references in academic research and teaching require the explicit approval of the editors and must state the authors and original source.

Experience Effects in Economics

Lessons from Past and Current Crises

Ulrike Malmendier

Presentation at the Annual Conference of the Priority Program

“Experience and Expectation. Historical Foundations of Economic Behavior,” June 2021

Keywords: lived-experience, experience effects, pandemics, expectations, beliefs

Let us consider the Covid-19 pandemic. How has life in quarantine affected us, and how will it continue to affect us? While we were living through stay-at-home orders, a lot of the discussion in journals and public media, but also among economists focused on the immediate impact of staying at home: For instance, what is the impact on consumption-behavior if I have less interaction at work or learn to do remote shopping? Also, as the *Wallstreet Journal* declared, everybody is a day-trader now. The other part of the discussion, which went a bit further ahead, centers on the consequences for the median income. There, economists often worry about the effects of the pandemic on earnings, wealth, job losses and job uncertainty – of yourself and of your children who might be moving back in at home. How does that impact educational choices, job choices etc.?

There was less discussion on the long-run impact: How does the experience of living through the Corona crisis alter our beliefs and behavior in the long run? And how do these effects depend, for example, on our personal exposure to the crisis? Let me illustrate this with a thought exercise. Imagine you are living in a country where large parts of the population are vaccinated, and the vaccine is effective against all variants of the virus. Everybody in this country has their jobs and job security back, their earnings and earnings prospects are the same and the impact on the cumulative wealth is minimal. Basically, it would be a pre-Covid-19 world again. My question is: Under these assumptions, would we also be back to economic decision-making and financial risk-taking as before?

If we are just focusing on the short- and medium-run impact, the answer would theoretically be ‘yes’. We have re-created the pre-Covid-19 world, and hence we should be making the same types of decisions. What I mean to convey here is that arguments about changes in the long run often come from the underlying assumption that economic conditions have actually changed. Of course, we are not exactly in the pre-Covid-19 world: industries have been disrupted, work processes have been changed, etc. But what about *us* having changed, *us* behaving differently, even if we were to return to that pre-Covid-19 world? That is my question here.

Previous epidemics or pandemics like the Bourbonic plague, tuberculosis or the 1918 influenza are generally recognized to have had a great impact on historical outcomes: on powers, wars, religion, economics, culture, etc. For instance, the Black Death clearly had an enormous demographic impact on Europe between the 14th and 18th century. Furthermore, there were changes in GDP, changes in trade patterns and routes, changes in financial capital available. All these sorts of outcomes are generally recognized in historical accounts. But there is another

aspect which may not be immediately apparent to economists, namely changes in world views and beliefs. For example, during the Black Death, there were changes in medical beliefs: People perceived the disease to be in the air and sought a solution in smoking. Much unlike the evidence we have today, they thought that this might help them destroy the illness. Also, people started thinking that maybe “it” attaches itself less to certain surfaces, which is why people started wearing leather-robos. But these are not the beliefs I have in mind here.

My research focuses rather on changes in beliefs about how the world is governed, in world views and religious beliefs. For example, the Black Death played an essential role in stimulating the development of a new piety, with counts of plague saints emerging. Also, modern theodicy discussions arose: How could God let that happen? But Yale-researcher van Snowden talks about how the result was not so much that some people became very pietist and some atheistic. It was more about, as he puts it, that there was a new kind of despair, a psychological shock that in historical hindsight, and with an anachronism, we might call “posttraumatic stress.” That brings me closer to the type of analysis I have done with more modern economic shock data.

One type of belief-change or worldview-change that also emerged from the time of the Black Death is the emphasis on the vanitas idea, the idea that earthly life is fleeting. It therefore might not be worth investing much, even in human capital – something I am currently very concerned about as it pertains to the current crisis. I started looking into this issue in the context of the very similar pandemic we had a hundred years ago, the 1918 so-called Spanish flu, the influenza pandemic. I was especially interested in how kids unable to attend school may have been harmed by this in the long run. One relevant variable about educational attainment from a hundred years ago is the high school completion rate. At the time – citing here US data – less than 20 percent of the population graduated from high school. But, starting with the birth years around 1880, it kept increasing, to 22, 23 percent, etc. Then the pandemic hit. The first generation affected by that were born in 1900, and so in school at the time and possibly deciding to change their career path as a result. Indeed, graduation rates experienced a sharp decline, not recovering for multiple generations. Schools were eventually reopened, making it possible for students to get their degree. Many, however, never returned. The trend of increasing educational attainment was disrupted and there was a downward shift in graduation rates for as long as two, three or even four generations.

What is worrisome is that this downward shift tended to be particularly strong for children from weaker socioeconomic educational attainment backgrounds. African Americans, for instance, had very low high school graduation rates of about 5 percent. There was an upward trend just

before the pandemic, which was quickly disrupted afterwards. You get an even better picture if you plot according to the educational achievement of the parents, socioeconomic status, or prestige of the parents' job, in particular the father's job. We see that for people who were high on the economic strata, two to four generations of kids were disadvantaged, after which they caught up to the trend on which they were before. For families on the lower end of the socioeconomic spectrum, however, the disruption endured for much longer.

That brings me to the general area in which I am interested, namely how traditional economic models of decision making, and especially belief-based decision making, leave out an important aspect by not considering the difference between learned knowledge description versus personally experienced outcomes. In a neoclassical model, the effect of personal experience of pandemics or crises would be no different from information about these outcomes, *ceteris paribus*. Living through a depression in financial investment should then be no different than reading about it. Having personally experienced unemployment should not affect one's consumption patterns any more than knowing about the risk of future unemployment. And living through a pandemic should be no different than knowing of its likelihood.

Emerging models and empirical evidence on the effects of experiences question that. Researchers like me and many others argue that personal experience has indeed a lasting impact on behavior. These impacts can occur on both the positive and negative side, but since I tend to specialize in crises and bad outcomes, I consider mostly scarring effects. Going slightly beyond my original field of expertise, I want to point to neuroscience and the fruitful field of research on the rewiring that goes on in the brain as people live through these experiences. In particular, processes of synaptic tagging and long-term potentiation can be very insightful to study.

To demonstrate that these effects are real and important, let me suggest another case: Depression Babies. People in the U.S. who lived through the Great Depression are known as the Depression Babies. The general view is that they, as a result of that experience, became very risk averse and avoided things like the stock markets or even flying. This is indeed the case. When starting our research, Stefan Nagel and I were surprised not to find anything on this in the finance literature, because it is so visible already in the raw data. If you look at stock market participation rates between the mid-30s and mid-40s and group people by cohort – so the one cohort those born up to 1920, the other from 1921 to 1930, the next from 1931 to 1940, etc. – you see that the generation that experienced the Great Depression as teenagers and young adults have a participation rate that is less than half that of all other cohorts. Of course, there is a secular upward trend, but this is a different type of jump. Also, there is future variation that

goes beyond just the Great Depression. For example, the thirty-one to forty cohort – people that experienced the post-war boom during their young adult life – has a participation rate a bit more than twice as high as that of the Depression Babies. Then there is a dip again for the forty-one to fifty cohort which experienced the stagflation years of the 1970s.

In our research project, we asked if we could predict risk-taking according to the domain in which people have experienced these bad outcomes. Here, the probability of individual i at time t of participating in the stock market. I try to predict this with all the usual variables we use in finance like demographics, individual wealth, etc. But then I throw in, in addition, some measure of prior lifetime experience. This is a little different to the lab environment, as these are decades of possibly accumulated experience. We take a weighted average of the most relevant outcome variables for experience – in this case, stock returns. The question is: “Do I get predictive power on stock-market participation?” So, we can take stock market experience, stock market returns or stock market participation; we can go to other financial markets, for example inflation experience related to bond market outcomes. And what we found is that the effects were very sizeable, statistically and economically highly significant.

By doing an inventory decile range comparison, I compare a person whose average life-time experience has been in the tenth percentile – bad stock-market outcomes in her life so far – to a person at the ninetieth percentile: That predicts a difference in stock market participation rates of forty percentage points. This is quite large compared to the average over the sample period, which is much of the twentieth century, with an average of 36 percent. There are similar differences for bond market participation. Here, too, the average participation rate for those periods is around 36 to 37 percent, whilst the inter-decile range of personal experience predicts an increase of 15 percent.

Interestingly, we did not find cross-fertilization. This means that experiences in one risky asset would affect the belief about the outcome of some other asset – at least from a standard financial model perspective and in the covariant matrix of how these outcomes relate to each other. But there seems to be no impact. So, experiences matter, they matter in the long run, and they tend to be very specific to the arena, the market, the environment, or the type of outcome that was experienced. With these results I refer to the domain notion that Elke Weber and her co-authors have used. But I am not entirely sure whether “domain” is the right word. There is a set of stimuli which triggers an experience-affecting behavior in those environments. But they are not one-to-one with the variables as economists code them up – that is what I want to point out.

How do people weigh their lifetime experiences? The data explain just that. We have looked at many different outcome variables from different markets, consumption, and unemployment experiences, etc. And we always wonder whether it is the case that early impressions leave a large impact, and therefore whether there should be extra weight on what is happening early in a person's life, whether the coefficients are completely flat and equally weighted, or if there ought to be decreasing weights. The latter is typically what we find. So, we model an A -function in the Probit model on experience effects as the weighted average of stock market returns starting today, then going back to $t-k$, and then all the way back to, say, the birthyear. I model the weight as a function of $H-k$ over a normalizing sum. That means, that if I put in a λ that is equal to zero, I would be getting equal weighting, λ equal to one gives declining weights. This means here that there is a function for a hypothetical 50-year-old person starting from today going back to the birthyear. There are linear declining weights. As λ goes up, it is getting steeper and steeper, and the other way around.

What the data would tell us in many different settings is that there is a recency bias, so there is more weight on what happened recently. However, even things far back in your past can leave a lasting mark, especially if it was a big shock such as the Great Depression. What you can do with this from a market perspective is quite interesting. You can aggregate the experience effects and ask whether we can explain the way markets move using this variable. Preliminary results suggest that this is the case. If we believe that individual-level estimates of past stock-market realizations affect a person's attitude towards the stock market today, we can aggregate them up and average them over the whole US population. So, I first go to the census and take every person between, say, 25 and 75 and take their lifetime experiences with the stock market – basically how the S&P 500 developed over their lives. Then I use a linear declining weighting function and come up with averaged experienced lifetime returns of the whole US-population. I plot against that the price-earnings ratio. They are highly correlated. So, at times when a lot of people in the population have had good experiences, society should be enthusiastic about being in the stock market. We see that the valuation of assets relative to the fundamentals is relatively high, and vice versa if lifetime-experiences go down. That is not mechanistic, it did not need to come out of the micro-level estimates.

You can also look at who is actually driving these changes in the stock market. In particular, the strength here is that we are not just fitting macro-data to some representative consumer and finding recency bias, adaptive learning, or similar results. We are saying that it matters how long your life has been. If, say, inflation in the U.S. peaked in 1980 and a person was born after that, she would probably be less reminded of that in current debates about possibly increasing

inflation. But if a person was born before, it may play a role. In other words, our differences in opinion may depend on our differences in lifetime experiences in the stock market, in inflation and all these variables we have looked at. Let us consider the example of the stock market. If you look at the difference in lifetime experience between older and younger cohorts – say, above 60 and below 40 – on the horizontal axis and plot that against the difference in stock market participation, you see a positive correlation. So, we can detect that not only does the average of the population matter, but we might be interested if it is the older or younger generation that is going in and out of the market. And that is strongly predicted by their lifetime experiences. So, while we often focus on age-specific determinants – “it is the older generation that invests more in bonds,” and so on – I maintain that we should not focus solely on age and life cycle considerations, but on the lifetime experience of that generation at this point in time. Sometimes it is the older generation, sometimes it is the younger generation that is more present in the stock market, for example.

If you embed this notion of how lifetime experiences affect your views of the world, your beliefs about future outcomes, in a simple OLG-model, as I did with my colleagues Demian Pouzo and Victoria Vanasco, you can generate the results I have shown so far. So, let’s say you have an OLG-model of finitely lived agents and every time the personal experiences they have – the dividend paid out of the Lucas tree – get extra weight, then you get belief-heterogeneity due to different histories. You also get an interesting implication, which is much neglected in the macro-finance literature: You get younger cohorts to react much more strongly to a dividend-shock than older cohorts, because the recent shock makes up a larger part of their lives. Again, going away from just age, a younger cohort reacts more strongly if a shock happens, while the older cohort has seen other things in their lives, averages it, and does not react as strongly. You generate these implications from market compositions I just showed empirically. So, if a big negative shock happens, the young cohort shies away and drops out of the market, resulting in a big overrepresentation of older cohorts.

But you can generate more predictions. For example, as the level of disagreement changes between cohorts, trading volumes should go up. However, since this has not been tested yet, we calculate a proxy for changes in experience-based beliefs. First, we have standard deviation in the experience of different cohorts – not just old versus young, but all the birthyears present: population weighted. Then, we plot against that a detrended proxy for trade volumes. This is the deviation of the turnover ratio from trend. It is linearly detrended and CF-filtered, so you have to adjust the trade volumes data because we know that there is a strong secular trend of it increasing; but amazingly it is closely related once you take out those additional tendencies.

And I think that there is more work to be done here by people more adept at writing down these asset pricing models.

Another thing to be done, maybe from the aggregate perspective, is to go to international capital flows. There are various international macro-puzzles regarding capital flows and portfolio investment. Scholars are still debating whether they have understood all the relevant explanations. In particular, with the same cohorts that we have looked at, there is the home-bias in equity: people holding an over-proportional fraction of their equity wealth in domestic stocks. The puzzle is that in periods of domestic crises (nobody likes that, whether you are a foreign investor or a domestic one), it is the domestic people who do not sell, and then end up being over-proportionately represented in that domestic market. The flipside of that is so-called “fickleness”, as Caballero and Simsek would call it. That is, if there is a foreign or global crisis, there is a withdrawal from foreign equity markets.

Our basic intuition of how experience effects can contribute, is to think about domestic versus foreign investors a little bit like the older versus the younger generation. The domestic investors have seen a lot of things. They have seen the domestic market for a long time, and experienced it for a long time, their lifetime, while foreign investors started coming in and personally experienced it only for a shorter period. So, you can think about it as these domestic investors having a more precise prior. As a result, they are not as shocked when a domestic crisis happens as the foreign investors, who like the younger people strongly update their beliefs when a crisis happens. So, the experience effect can explain all three of these core puzzles on international capital flows and investment.

It can also then generate additional predictions that could be tested. For example, the strength of these puzzles should depend on the demographic composition in different countries. It should also depend on whether we are looking at this in times of generally higher or lower economic activity. For example, in a country with a larger number of young market participants you would expect people to overreact more than in a more ageing society. And they would overreact both to domestic and foreign shocks – because they have neither much experience with their own market, nor with the foreign market. As a result, they act more like these foreign investors, who strongly overreact to a domestic crisis. The whole retrenchment- and fickleness-puzzle should be alleviated, because everybody is withdrawing, not just the foreign investors over-proportionately. Using data from the IMF, World Bank and World Federation of Exchanges, we were able to confirm this. Still, I am not an international macro or finance person, and I

believe much more could be done. There are a lot of low hanging fruits in this international context.

Now, in terms of identifying variations of experience effects, what I have been using so far is the birth year. Cohorts are differently exposed to macro-level realizations. This links us to thinking about the generational effects of, for example, the Gen-X or Gen-Y or Millennials being different from other cohorts. Right now, a lot of us are probably asking “Who will be Gen-Covid? Along what dimensions will they be scarred?” – so, looking over time at differential exposure to different macro-finance outcomes. I think this is a very useful tool to detect experience in field data.

However, you could also clearly think about the exposure varying by location – think about Covid, some countries and regions have been much more strongly affected than others. And then, even within regions and within periods of time there are individual differences that play a role, which in some sense are what experimentally are exploited if we give people different treatment. One example to show these different sources of identification is what Lesley Sheng and I used in our paper on scarred consumption, where we try to show how, in order to predict one's current willingness to spend – their consumption expenditure – we should go beyond the usual demographic-, wealth-, and income-controls and beyond other life-cycle considerations and also consider what a person has lived through in her life cycle so far.

Consider, for instance, two people with the same income, same income prospects, same wealth, employment prospects etc., but one person made it from a poor background where she was struggling, had to work herself up and thus knows what it is like to not have a lot of money to work with, while the other person was born with the silver spoon in their mouth. I would then predict that they make different consumption expenditure decisions, with the first person being much more careful than the second one. To test this, we related consumption spending in the PSID to peoples' individual unemployment experience – but also to what they have seen around them, whether locally or nationally in terms of unemployment, people losing their jobs, etc. – in a regression that controls for time fixed-effects, state fixed-effects and even household fixed effects. That is the nice thing about using panel data.

I want to illustrate how you can identify experience effects in such data with three handpicked people from the PSID data. Here, I have a person born in 1948 that lives in Pennsylvania. That person has, as of 2007, before the financial crisis, witnessed an average statewide unemployment rate of 5.81 percent. That is lifetime-experience with linearly declining weights. Now, there is a person of the same age, born in 1948, but living in Alabama. Average

unemployment was somewhat lower here, 5.7 percent. Let me compare those two people first. The first person is spending a little bit more than the second one – to nicely fit my story, I picked them that way. Now, the financial crisis begins, and it hits Alabama stronger than Pennsylvania. Six years after the crisis, in 2013, the lifetime experience of the Pennsylvania person is now a 6.06 percent unemployment rate, while the Alabama person, who had a better experience so far, now has a lifetime-experience of 6.11 percent. Correspondingly, we see in the data that the Alabama person is reducing her spending much more drastically than the person from Pennsylvania – by 13 and 7 percent respectively. And, just for completeness, to say that the time dimension is still there: Let me add a third person here, also living in Alabama, but younger. Born in 1975, that person is 32 years old when the crisis starts. Their whole lifetime experience has been very good – 5.46 percent average unemployment. But the crisis means that her lifetime average drastically worsens. Their lifetime experience of unemployment jumps from the lowest to the highest, 6.2 percent. Now, that person is spending less to begin with – that person is younger, so life-cycle considerations clearly play a role here. But they then reduce spending by 21 percent. This shows how in addition to the time dimension, using geographical variation, individual differences and experience parallels can play a significant role and allow you to detect experience effects. Overall, we found in the data that personal experiences and also macro-experiences or local experiences significantly predict lower consumption spending as unemployment goes up. They also predict more pessimistic views about financial conditions.

The beauty is that they do not predict lower personal income or more volatile personal income after including the controls I have been using. These people live through unemployment around them or personally, become pessimistic and very cautious spenders, even though they have no reason to do that. I am controlling for their income, income prospects, etc. Instead of consuming, these people build up higher wealth. They start to save. There is a really strong impact on saving that can be explained both by cohort-specific variation and by local variation.

To make a brief intermediate summary: We have seen in some of the evidence that experiences over one's lifetime have a long-lasting effect on beliefs and economic choices. Different cohorts are affected differently and their relative position changes over time, and it is not an age effect. We have seen that experiences are domain specific – there is no cross-fertilization between realms of economic decisions. We have seen that the extent of the exposure matters, and that different locations can be affected differently. That, just as an aside, has important implications. Think about the Covid-19 crisis and how different genders were affected differently, depending, for instance, on care responsibilities. Different races have likewise been affected differently, with minorities like Latinos and Blacks being over-proportionately affected by the current

crisis. We therefore predict that they are altered differently in their choices – investment choices, job choices, saving and consumption choices – in the long run and even after this exposure has passed. This has direct implications for inequality.

The last point I want to add in terms of stylized features emerging from this literature is that these experience effects seem to play a role even for experts. In behavioral finance in economics, we say that there is a bias that may affect the individual consumer and individual investor, but surely not the professional. But I argue that these experience effects matter for everybody. One example is inflation. I mentioned briefly that lifetime exposure to high inflation has strong predictive power in tilting beliefs and worries about future inflation. And that is the case not only for individuals or experts. To pick a German motivation: You can think about the long-lasting effects of the German hyperinflation, either personally or quasi-personally, growing up with pictures in textbooks and the stories of grandparents. But also in the U.S., when reading about the 1970s high inflation, peaking in 1980, it is very easy to find statements by then-chairman Paul Volcker, for example, lamenting that a young generation of adults has only seen increasing inflation: How can you possibly convince them that they are going to go back to price stability? Volcker knew they had experienced just this, and that you could teach them whatever, give them descriptions however much you wanted to, but that they nevertheless might not change their worldview. In a previous paper with Stefan Nagel, we said that when people form inflation expectations, they put more weight on realizations they have personally experienced over their life. We discovered a weighting pattern which is again very similar to the stock market with linearly declining weights and showed that it can have a significant impact on financial decisions and borrowing and lending.

Now, let's go to the professionals, namely to the members of the Federal Open Market Committee, the core central bankers in the U.S. In particular, consider a man who was born as Heinrich Wallich in Berlin in 1914 into a family of bankers that lived through Germany's hyperinflation and then emigrated to the US in the 1930s, where he had a successful career as an economist and FED-governor from 1974 to 1986. Heinrich Wallich still holds the record of dissenting 27 times against chair proposals of setting the rate, where people might agree or dissent into hawkish directions, saying that higher rates were needed to lower inflation; or that lower rates were needed to combat unemployment. Wallich still referred to his personal experience, insisting that people did not grasp the dangers of inflation decades later. He was highly educated, knew it was a different monetary system, different country, different decade etc. Wallich is my favorite example, but we can generalize: If you plot FOMC members' inflation forecasts, normalized by staff forecast, against their personal lifetime experience with

inflation, you indeed find a really strong positive correlation. So, everybody seems to be affected by what they have seen throughout their life – and again, these are *the* central bankers of the U.S.

That gets me to the channel, or underlying mechanism, of this discussion. When you present this kind of experience effect evidence, a lot of economists say that these are informational frictions. People do not have access to information, but by personally experiencing something they access it. I hope that with this evidence on central bankers and other professionals – I am also working, for example, with doctors – I can rule this out. Scholars with a more traditional behavioral finance background, and a behavioral economics approach, often argue that there are some biases and that we are not processing this information correctly. Here too: The FOMC members have all the inflation models and all the possible information available. So, are they really not able to put in the data and come to the right conclusion?

If it is not a bias in the traditional behavioral sense, which I was calling “software-failure” here, maybe we have to go to the hardware. Maybe something changes in our hardware as we are having this experience. In particular, the concept of synaptic tagging has high potential. These are new models of belief formation which are emerging now in macro-finance. The idea is that every time we live through a new experience, we have some neurons firing, some presynaptic neuron telling some postsynaptic neuron whether to get agitated or to calm down, etc. Those synapses form the way we experience life around us. Not only children, as we used to think, but even adults have brain plasticity throughout their lives and the brain never stops changing as we are living through experiences. In particular, neuroscientists tell us that how and how often we have an experience matters. If the hippocampal neurons are repeatedly stimulated that can induce increased strength. There is thus a long-term potentiation which triggers the brain when the stimulus comes in again to immediately go to that fear-reaction or positive reaction that was so strongly strengthened during previous experiences. Knowledge has very limited power to undo these effects.

One last thing I want to add here: I very much like to think about trauma and traumatic stress causing these synaptic changes as evidence that can help us better understand belief formation in economics. Some of my colleagues tend to call my research interests an interest in econ-PTSD. And I quite like that. The economic crises we live through lead to some re-wiring, which leads to post-traumatic changes in our behavior. Now, the only issue with that is that people then tend to only think about the very big crises – and of course I reinforced that by talking about hyperinflation, Covid-19, the Great Depression, etc. But, as neuropsychiatrists will never

tire to underline, there is not only Trauma with a big “T” like war-experiences, accidents, child-abuse, adverse childhood experiences – or in econ, the German hyperinflation, the Great Depression, pandemics. There is also a trauma with a small “t”: the daily micro-aggression, the daily swipe, the daily paper cuts that get reinforced in the brain and let people think about themselves and the world in a different way than they would have had they never been exposed to them. These experiences can have as strong effects as big traumatic events. And, in fact, they can be much harder to combat because it is hard to identify them and to get rid of them. I think that this translates into economics. I think the, say, daily worry about food, prices, or unemployment can re-shape your brain and have a long-lasting effect on decision making.

So, to give one example of that, let me refer to a paper of mine on gender roles and the gender expectation gap that recently came out in PNAS with Francesco D’Acunto and Michael Weber. In this paper, we build on the long-known puzzle that females tend to be much more pessimistic about future inflation than males. In the FED, everybody is aware that women have much more pessimistic, higher inflation expectations than males. We replicate that by regressing inflation expectations on gender, and even with a lot of controls and in periods of low inflation like that which we had until recently, and which we used for the measure, there is an economically significant half percentage-point difference. However, we then zoomed in and added a survey to the NIELSEN-data that elicited beliefs and underlying shopping behavior. We asked the question: “Who in your household does the grocery-shopping?” What we found is that the gender difference completely disappears if we look at the households where both males and females are participating in the grocery shopping. Instead, the puzzle is completely driven by those families where, in a more traditional gender role type of setting, the female is doing all the grocery shopping. How can that be? Well, grocery prices are very volatile – in fact, that is why they are taken out of the core-inflation index or CPI in the US: economists cannot detect very well any inflation trends, and we know from previous research that people anchor on the increases, while they do not put a lot of weight on the decreases. So, by giving a lot of weight to grocery price inflation, people see a lot of increases and think about strong inflation. That is how the puzzle emerges that women have a more positively biased inflation expectation, and that is how the observation completely disappears if we control for grocery-shopping.

Let me sum up. I wanted to give you a taste of the existing and ongoing research on experience effects in economics. The idea is that if you try to predict some kind of outcome variable Y , say spending, and you think you have all the relevant predictors – the vector X for person i at times t : Maybe think twice and go back to previous realizations of that vector. Even past unemployment experiences, or other past exposures, might have a significant outcome on

behavior, even if it does not have predictive power for future outcomes above and below the Xs you are already controlling for. We have evidence for that from finance and from macro, but there is also related evidence from labor economics, political economy, and other applications. And, in general, the effect we have been distilling is that there is a very long-lasting impact, combined with some recency bias, domain specificity and a neuro-science foundation that affects even experts – there seems to be some real neuro-rewiring.

In terms of future research, I think that there is a lot more to be done in the realm of household finance, on gender, on the long-lasting effects of exposure to racism, of exposure to adverse title experience, on education more generally. Luckily, we live in a time where big data even on that taking place “within” persons will make this possible. And I predict that this will have very important welfare and policy implications. And with that, I will stop here.