Research Statement

Qilin Zhang

I am a macroeconomist who studies topics in public and family economics. I use microdata and quantitative models to study how households and governments respond to institutions and public policy: why fewer people marry, how parents invest in their children, how governments finance public investment under fiscal constraints, and how families respond to local labor-market shocks. I examine these questions in China, where marriage and local public finance have changed rapidly, and in the United States.

Job market paper: gender-biased dynastic savings

My job market paper, "Gender-biased Dynastic Savings and the Marriage Market: Evidence from China," studies how parents' investments in their children shape the marriage market. China's crude marriage rate fell from 9.9 per 1,000 people in 2013 to 4.3 in 2024, and because births outside marriage are rare in East Asia, the decline in marriage drives the region's low fertility. Existing explanations look at one side of the market: why men need ever more assets to marry, or why highly educated women remain single. I argue that both are outcomes of the same parental behavior, which I call gender-biased dynastic savings: parents save for their children in two forms, education and financial assets that the child brings into marriage. The paper asks two questions. First, how does this investment mix differ by the child's gender? Parents tilt daughters toward education and sons toward assets. Second, does this gender-biased investment matter for the marriage market? It does: it lowers the marriage rate. In a dynastic model of parental investment and marriage matching estimated on household survey data, rising parental resources along China's growth path account for about 35 percent of the percentage decline in the crude marriage rate between 2010 and 2024.

Three facts from the China Family Panel Studies answer the first question. First, marriage sorts strongly on education, and single women are the most educated group. Second, education spending rises with family income more steeply for daughters than for sons: the income elasticity of educational expenditure is higher for daughters by 0.08 to 0.14 across samples, and the gap also appears within families. Third, parents save more for sons' marriage: households with sons have higher saving rates, and the financial-asset share of sons' parents falls by about 4 percentage points when the son marries.

Two gender asymmetries explain why parents split their savings differently by the child's gender. The technology that turns education spending into attainment differs by gender, and for daughters it varies with family resources. The marriage market sorts on education and expects husbands to bring assets. Each family's choice is individually optimal, but together the choices determine the education composition and relative asset positions of the next cohort, and with them how many compatible couples can form.

To quantify this mechanism, I build a model of one-child families that differ in parental resources and in the gender of the child. Altruistic parents allocate resources among their own consumption, education spending, and a premarital asset transfer. Children then enter a transferable-utility marriage market with idiosyncratic preferences, in which couples pool their assets, education enters the match surplus, and a husband-provider norm penalizes husbands whose premarital assets fall short of their wives'. I estimate the education technology and the marriage surplus from CFPS data and calibrate the remaining parameters to the gender gap in the education–income elasticity, the saving-rate gap between sons' and daughters' households, and selection into marriage. The estimates show that the education technology is more responsive to family resources for daughters than for sons, and the model matches these targeted moments.

Three counterfactuals answer the second question. First, equalizing the education technology across genders raises the marriage rate by 0.62 percentage points through two opposing channels. More sons become highly educated, which lowers the marriage rate by 0.44 points. But the sons who remain low-educated enter the market with more premarital assets, which the norm rewards, and this raises the marriage rate by 1.05 points. What

matters is therefore less how much education sons receive than which sons enter the market with few assets.

Second, economic growth amplifies this gender bias, because daughters' education spending responds more to family income. As parents grow richer along China's realized GDP path, the model's marriage rate, the share of the marriage-age population that is married, falls from 64.4 percent in 2010 to 52.3 percent in 2024 and to 43.3 percent by 2035; the 2010–2024 decline of 18.8 percent is the basis for the 35 percent figure above.

Third, transfers do not automatically raise marriage. A transfer equal to 20 percent of mean parental resources lowers the marriage rate by 8.73 percentage points when paid to parents, because it funds more education, and raises it by 14.22 points when paid to newlyweds. Where a policy lands in the family's life cycle matters as much as its size.

Public finance and local government debt

My second project turns from households to the governments whose policies they respond to. It starts from a basic public finance question: when formal fiscal resources are limited, how do local governments finance investment, and how much do they borrow outside the budget? In China, local governments carry heavy responsibilities for investment and development but have limited ordinary-budget revenue. My solo-authored paper, "A Micro Anatomy of the Hidden Local Government Debt in China: 2014–2023," presented at the 2025 Southern Economic Association Annual Meeting, builds a firm- and city-level dataset from public disclosures on 2,940 industrial local government financing vehicles (LGFVs) across 305 cities. LGFVs are state-owned firms that borrow on behalf of local governments outside the official budget. The paper documents wide heterogeneity across regions and sectors: hidden local government debt grew from RMB 10.2 trillion in 2014 to RMB 50.9 trillion in 2023, and inland cities rely more heavily on hidden borrowing. The dataset supports future work on the geography of off-balance-sheet debt, fiscal transparency, and debt sustainability.

Ongoing and future research

My next projects extend the same question in two directions: across generations, and to labor-market shocks in the United States. The counterfactuals in my job market paper hold fixed the marriage market that parents expect their children to face. The next step is to close the intergenerational loop: a steady-state analysis in which parents anticipate the marriage market their investments create. A further step is to let the sex ratio and the provider norm, which the model holds fixed, evolve as women's education rises. Together these would show whether gender-biased dynastic savings are self-reinforcing across generations.

Turning to the United States, in "Mass Layoffs and Fertility: Evidence from WARN-Based Shocks," joint with Yuchen Zhu, I study how local labor-market shocks affect fertility. Using WARN mass-layoff notices linked to county-level natality data from 2020 to 2024, we examine how the timing and intensity of layoffs relate to subsequent fertility. Preliminary results suggest that larger layoff shocks are followed by higher fertility, with the response emerging about ten months after the shock and driven mainly by large layoffs. Ongoing work examines the roles of the opportunity cost of time and maternal selection. In "Residential Mobility and Household Consumption," joint with Yanqi Li, I use PSID data to study how household consumption changes around a residential move.

Across these projects, I see economic research as detective work that ends in a story. I start from facts in detailed microdata, often descriptive, that point to a question worth asking and discipline the assumptions of a model, and I build the simplest model that can tell the story those facts suggest. I use causal designs when identification is the main question and quantitative models when the question is about mechanisms and counterfactual policy, aiming for technical rigor in service of an important question. Rather than ask readers to take my conclusions on trust, I try to lay out the evidence and the reasoning clearly enough that they reach those conclusions themselves.