Soft Power Shocks: Does Attention to a Symbol Translate into Bilateral Exchange?

Zong-Wei Yeh*

This version: November 2026

Abstract

Governments and zoos credit the giant pandas that China lends abroad with more trade, more Chinese visitors and warmer feelings toward China. We test this expectation using a soft-power shock to attention to China's symbol. The shock comes from cub births at foreign zoos, whose timing is set by biology and involves no decision by either government. We measure each birth's dose in the host public's Google searches. A birth raises attention to the panda, but the attention does not convert. Searches for China, Chinese visitor arrivals, imports of Chinese consumer goods and bilateral trade show no detectable response. Trade effects of more than about 3 percent per typical birth are ruled out, against estimates of 5 to 7 percent. Political conflicts with China cut Chinese visitors and the targeted exports. Attention to a symbol stays with the symbol, while conflicts with the Chinese state move the same flows.

Keywords: Soft power, panda diplomacy, attention shocks, local projections, bilateral trade
JEL codes: F14, F51, E71, L83

*Corresponding Author: Zong-Wei Yeh, National Chung Cheng University, Department of Finance, No. 168, Sec. 1, University Road, Minhsiung, Chiayi 621301, Taiwan, E-mail: zweyeh@ccu.edu.tw.

"Pandas have long been envoys of friendship between the Chinese and American peoples."

Xi Jinping, speech at the welcome dinner by friendly organizations in the United States, San Francisco, November 15, 2023 (Xi 2023)

1 Introduction

Governments, zoos and the press expect the giant pandas that China lends abroad to bring visitors, trade and goodwill. When two pandas arrived at Edinburgh Zoo in December 2011, an assessment prepared for Scottish Enterprise, Scotland's economic development agency, saw in them a chance for the city to "punch above its weight" in attracting visitors and building trade links with China, and recommended branding Edinburgh "as a 'China friendly' place to visit and do business."[1]

The expectation is widely shared, and it has been quantified. Exports to China are found to rise by 5 to 7 percent in the year a cub is born at a foreign zoo (Longhi and Morelli 2025).[2] Countries that host pandas are found to receive more Chinese tourists, with the effect strongest after three years of hosting (Okafor et al. 2021). The economic effect of one cub at Ueno Zoo in Tokyo was forecast, before the cub was born, at 26.7 billion yen a year (Kansai University 2017), and Adelaide's pandas were expected to bring A$632 million to South Australia over ten years (Australasian Leisure Management 2009). Beyond the numbers, young people abroad are expected to "develop . . . an interest in various aspects of Chinese culture" through the pandas (Liu 2024), and China's forestry administration states that its panda cooperation promotes friendly exchanges between the Chinese and foreign publics.[3] Together, these statements describe an expectation that public affection

[1] Scottish Enterprise (2012), pp. 2–3 and 20. The assessment estimates the additional visitor spending in Edinburgh over the ten years of the loan at between £2.8 million and £27.6 million, depending on how actively the city and its partners promote the pandas.

[2] Longhi and Morelli (2025) estimate a gravity model by Poisson pseudo-maximum likelihood on annual Comtrade data for 2002–2019 and find exports to China 5.3 percent higher in the year of a birth across all countries and 7.4 percent higher among panda hosts, with the birth year running from November to October. They also find that a monthly index of political relations with China improves after births, with the effect concentrated in hosts less aligned with China.

[3] China Daily (2023) reports the National Forestry and Grassland Administration's assessment of its

for China's symbol carries over to attention to China, to Chinese visitors and to trade. This paper provides a test of that expectation with an identified shock.

Whether affection converts is a question about the channels through which relations between countries reach economic exchange, and two channels can operate at the same time. Public attitudes toward a country are associated with how much its partners trade with it and invest in it, as consumers, travelers and investors favor the countries they regard well (Guiso et al. 2009; Disdier and Mayer 2007; Rose 2016). Political relations also move exchange through the decisions of governments, through state-owned importers, public procurement and the rewards and punishments that governments hand out (Fuchs and Klann 2013; Davis et al. 2019). The two channels usually move together. The question for cultural diplomacy is whether the attention it creates in a foreign public produces economic exchange when no change in policy accompanies it.

Answering the question raises three identification challenges. First, sentiment responds to economic and political relations, so that its movements are endogenous to the flows it is meant to explain. Second, the symbolic gestures that governments make are bundled with state action. China announces the loan of a pair of pandas during state visits, together with trade and investment agreements, which is why the panda loan itself cannot serve as a sentiment shock. Third, public attention to a symbol is not sentiment toward the country it stands for, and a measure of the first may move without the second.

We address these challenges with the births of giant pandas that China lends to foreign zoos. The births identify a soft-power shock (SPS), an exogenous rise in a foreign public's exposure to a country's soft-power asset that carries no change in policy. Panda diplomacy is the policy, and goodwill toward China is its stated aim. We hold the policy fixed, since the pandas are already on loan, and identify shocks to the public exposure through which the policy is meant to generate that goodwill. The paper therefore tests the premise of the

cooperation with institutions in 19 countries. According to the assessment, international panda cooperation has played an important role in "促进中外民间友好交往" (promoting friendly exchanges between the Chinese and foreign publics).

policy rather than the loan decision, which comes bundled with state visits and agreements. China leases pandas to zoos in 24 jurisdictions, and the births of cubs there draw national attention in the host countries. Their timing is set by the biology of the species. The embryo implants after a delay of variable length, and a female's hormones and behavior after ovulation look alike whether or not she is pregnant, so that a birth is hard to foresee even for the zoo (Zhang et al. 2009; Willis et al. 2011). On the day a cub is born, no government in Beijing or in the host country takes a decision. We assemble a database of 491 dated events in the lives of the pandas in the 24 host jurisdictions from 1980 to 2026, with the date of the first public announcement of each, and measure the size of each event by how much it raises the host public's Google searches for the giant panda in the days after the announcement relative to the days before. The doses of the births, summed by country and month, form the soft-power shock, which serves as our instrument for the host public's attention to the panda, an application of the narrative event-instrument approach of Romer and Romer (2004), Känzig (2021) and Gavriilidis et al. (2026).

The births move attention to the panda and not to China. The announcement of a birth raises attention to the panda to about nine standard deviations above its pre-announcement level on the day of the announcement, after which attention settles within a week to a level that stays above the pre-birth baseline for three months. Attention to China in the same countries does not move. The births are unrelated to geopolitical risk, policy uncertainty and oil and carbon policy shocks, and lagged attention, macroeconomic and political series do not predict them. The first stage is strong, and it holds when the dose is measured at the zoo's region rather than nationally, when the largest events are removed, and when each event is counted without its dose.

We follow the attention a birth creates up a ladder of outcomes, arranged from attention to flows, and find that it stops on the first rung. In the days after the announcement, the host public's searches for Chinese brands, travel to China, the Chinese language, Chinese culture and China itself move by a small fraction of the rise in searches for the panda, in the

same search engine and over the same days, and none of the responses is distinguishable from zero. The slower rises in searches for Mandarin and for Chinese destinations in the following month are calendar effects of a birth season that runs from July to September, and they appear in the same months of years without a birth. The attention a birth creates attaches to the symbol and does not pass to the country that lent it.

Nor do the flows named in the expectation respond. Chinese visitor arrivals in the host do not respond at any horizon. Host exports to China and imports from China are precise zeros. Effects larger than about 2 to 3 percent of the flow per typical birth are ruled out, well below the 5 to 7 percent of the gravity estimates. Imports of Chinese consumer goods, the category in which affection for China would show first, are zero with the same precision, and so is China's share of the host's consumer imports. These results hold across the dose-weighted versions of the shock, fixed effects, control sets and samples.

The same outcomes respond to conflicts with China. A conflict with China, such as a head of government's meeting with the Dalai Lama, a sanction, a territorial or Taiwan-related dispute or an arrest, lowers Chinese visitor arrivals in the host by about 17 percent within six months, in the same equation in which a birth produces no detectable change. China's sanction episodes cut the host's exports to China in the targeted product chapters by tens to hundreds of log points, while summits leave the outcomes unchanged. These events are chosen by governments, and they serve here as a comparison.

The paper makes three contributions. First, it identifies the SPS, a dated shift in a foreign public's attention to a country's symbol, measured in its dose and free of state action. Second, it tests a widely shared expectation about the economic returns to panda diplomacy with estimates whose precision is stated in the units of that expectation. Third, it compares the response of the same flows to conflicts with China in the same panel.

Related literature and contribution. Our paper contributes to three strands of literature. The first studies whether political relations between countries move economic exchange.

Nye (2004) defines soft power as the ability to obtain what one wants through attraction rather than coercion or payment. Public attitudes are associated with exchange. Bilateral trust predicts trade and investment between countries (Guiso et al. 2009), bilateral opinions predict trade (Disdier and Mayer 2007), and countries perceived to have a positive influence in the world export more (Rose 2016). Studies that isolate political shocks have mostly examined the opposite of attraction. Countries whose leaders officially receive the Dalai Lama see their exports to China fall, mainly in machinery and transport equipment, and the effect disappears in the second year after the meeting (Fuchs and Klann 2013). Imports controlled by state-owned enterprises respond more to political relations than imports controlled by private firms (Davis et al. 2019). Boycotts, nationalist episodes and other political shocks reduce trade and firm values, mostly briefly (Fisman et al. 2014; Heilmann 2016; Du et al. 2017). In all of these settings the shock is a conflict between governments, in which state action and public anger move together. We provide a shift in public attention that comes with no state action and estimate it in the same equations as the conflicts, so that the attraction side of soft power and the coercion side are measured on the same outcomes.

The second strand studies panda diplomacy itself. Panda loans are an instrument of China's public diplomacy (Hartig 2013), and since the 2008 Sichuan earthquake loan deals have been associated with countries that supply China with resources and technology (Buckingham et al. 2013). Two recent papers estimate economic effects. Longhi and Morelli (2025) use the timing of cub births as a shock to political relations in an annual gravity model and find that exports to China rise and that a monthly index of political relations improves. Okafor et al. (2021) compare Chinese tourist flows to countries with and without pandas and find more visitors to hosts. We measure the dose of each birth in the host public's attention, follow it month by month with country fixed effects and pre-announcement controls, and show in Section 4.5 how the two designs differ and why they give different answers.

The third strand identifies macroeconomic shocks from narrative events. Narrative

identification (Romer and Romer 2004) and event instruments scored by high-frequency market prices or news coverage identify oil supply news, carbon policy and climate policy uncertainty (Känzig 2021, 2026; Gavriilidis et al. 2026), and non-fundamental sentiment shocks can be identified from events unrelated to the economy (Lagerborg et al. 2023). Search volume measures attention directly (Da et al. 2011). We carry the narrative-event design, with local projections and external instruments (Jordà 2005; Stock and Watson 2018; Barnichon and Brownlees 2019), to a bilateral attention shock in a panel of host countries, with search data measuring the dose of each event.

The paper is organized as follows. Section 2 defines the soft-power shock and measures attention to the panda in each host country. Section 3 builds the soft-power shock, documents its diagnostics and sets out the econometric framework. Section 4 follows the SPS up the transmission ladder, from attention to the panda to attention to China, Chinese visitors and trade, and reconciles the results with the gravity evidence. Section 5 estimates the effects of conflicts with China in the same panel, and Section 6 concludes.

2 Measuring Attention to the Panda

Public sentiment toward a country is not observed at a monthly frequency. This section defines the soft-power shock and describes how we measure attention to the giant panda from Google search data.

2.1 Definition

We define a soft-power shock (SPS) as an exogenous change in a foreign public's attention to a symbol associated with another country, with policy held fixed. No tariff, agreement, sanction, state visit or government purchase accompanies it. The shock captures exposure, not the direction of sentiment. Whether the exposure produces interest in the country or exchange with it is the empirical question. The shock is bilateral, since it concerns the

public of one country and a symbol of one partner, and it is non-fundamental in the sense of Lagerborg et al. (2023). It changes no budget set or technology.

The giant panda is a natural carrier of such a shock, and the institutions of panda diplomacy show why. During the Mao era China gave pandas to friendly states as gifts, among them the pair sent to Washington after President Nixon's visit in 1972. After 1978 the gifts gave way to loans for a fee, and the agreements of the past three decades are long-term programs of cooperative research and breeding. They typically run for ten years and are renewable, with an annual fee paid to a Chinese conservation body, and cubs born abroad are sent to China after a few years (Buckingham et al. 2013; Smithsonian's National Zoo and Conservation Biology Institute 2025). China describes the purpose of the program as conservation, joint research and friendly exchanges between peoples, and calls the pandas envoys of friendship (China Daily 2023; Xi 2023). The hosts expect visitors, trade links and goodwill (Section 1). The panda is China's most distinctive instrument of public diplomacy (Hartig 2013), and its loans are negotiated between governments and often announced at summits (Section 3.1). A birth is a different kind of event. Its timing is set by biology rather than chosen by anyone, and its announcement is a dated event in one host country, covered there as national news. Appendix B describes the institutions and their history.

The definition has two implications for measurement. First, we exclude state action. Loan agreements, arrivals, returns and recalls are decided by governments and often announced during state visits. They carry soft-power content, but it comes bundled with policy. Our baseline therefore uses only biological events, and among them only those that did not coincide with a state visit or a diplomatic conflict. Second, because the measure is unsigned attention, Section 3 examines whether attention to China responds to births, and Section 4 follows the shock to bilateral flows. We measure the host side only, the attention of the host public to the panda and to China.

2.2 Measurement with Google search data

We measure attention with Google Trends. Rather than a keyword, we use Google's topic entity for the giant panda, which pools searches for the animal in any language and script, and the topic entity for China as a comparison.[4] For each of the 24 host jurisdictions we retrieve the monthly index from January 2004 to December 2025, and for each dated event a daily window from 60 days before to 120 days after the announcement.

Google scales each request so that its maximum equals 100, so a daily window and the monthly index are on different scales, and windows of the same country are not comparable with one another. We put every daily window on the scale of the country's monthly index by regressing the monthly index on the window's monthly means, through the origin and over the five or six months the window covers, and multiplying the daily values by the estimated slope. We then divide by the standard deviation of the rescaled daily series on non-event days, that is, outside the interval from 10 days before to 40 days after every dated event of the country, computed only on windows in which fewer than half of the days are zero. Denote by $a_{i,d}$ the resulting attention of country i on day d, in units of that country's standard deviation. A change of one unit in $a_{i,d}$ has the same meaning in Austria as in the United States. It is one standard deviation of that country's ordinary daily fluctuation in interest in the panda.[5]

For each event we also retrieve the daily series for the region of the host zoo where Google reports it. Regional series sharpen the dose for zoos far from the capital, but they are mostly zero outside large regions, and they are usable for fewer than half of the events. Our headline measure therefore uses the national series, and the regional series enters as a

[4] In South Korea, where Google is not the dominant search engine, the English keyword "panda" returns zero on most days around Fu Bao's birth, while the topic entity returns a value on all but two of 122 days. The topic identifiers are `/m/03bj1` (giant panda) and `/m/0d05w3` (China).

[5] All series are single pulls from Google Trends. In Spain and Mexico the panda topic is very high in 2004–2006 and then decays, most likely because searches for an antivirus company of the same name were mapped into the topic in those years. The contamination affects the level of the monthly index in those years but no event in our instrument, because the Spanish births are from 2010 onward and the Mexican births predate Google's data.

robustness version of the instrument.

Appendix Figure A.1 shows the monthly panda index for each host from 2004 to 2025. Most series share one shape, a stable baseline with spikes when pandas arrive or cubs are born. In larger hosts with several pairs, such as the United States, the spikes sit on a higher and more variable baseline. Japan's index jumps to a permanently higher level with the birth of Xiang Xiang at Ueno Zoo in June 2017, which reflects the coverage of Japanese-script searches by Google's topic entity rather than a change in interest, and births at the Wakayama zoo before then leave almost no trace in the national series. Section 3.4 states how we treat them.

3 Identifying Soft-Power Shocks

The three challenges set out in the introduction are endogeneity of sentiment to relations, bundling of symbolic gestures with state action, and the gap between attention to a symbol and sentiment toward a country. This section builds the soft-power shock, which addresses the first two, examines its validity, and sets out the econometric framework in which the third is tested.

3.1 Panda births as soft-power shocks

The SPS starts from a record of every dated event in the lives of the giant pandas on loan to the 24 host jurisdictions. The events are agreements, arrivals, public debuts, announced breeding attempts and pregnancies, births, failed pregnancies, namings, cub debuts and departures, extensions, returns, recalls and deaths. Table 1 classifies the 491 events by type, and Appendix Table A.1 by host. Each row carries the physical date of the event and the date of its first public announcement, which is the date we use, and nine in ten rows cite a primary source such as a zoo release, a wire story or a national broadcaster. Each event is also coded for bundling, with +1 if a state visit, summit or major agreement between

the host and China took place within seven days of the announcement, −1 if a diplomatic conflict did, and 0 otherwise. Bundling is common for the decisions of governments and rare for the events of biology. Of the 25 loan announcements, 18 were made during a state visit or summit, against one of the 67 births.[6]

Of all events, births come closest to random assignment in time. In giant pandas the fertilized embryo implants in the uterus only after a delay whose length varies across pregnancies (Zhang et al. 2009). Zoos report gestations ranging from about 95 to 160 days. After ovulation, females that are not pregnant go through a pseudo-pregnancy whose hormonal and behavioral signs resemble those of a pregnancy (Willis et al. 2011), so that many announced pregnancies end without a cub. The month of a birth is therefore hard to foresee for the public and often for the keepers. The season is not. Births cluster in late summer, which the month effects of our specification absorb, together with country-specific calendar-month effects where an outcome has its own seasonal pattern. Births also carry no decision by either government, and their first announcement comes from the zoo, usually on the day of the birth or the day after.

We measure the size of each event by the attention it draws. For each event e in country i announced on day d_e, the intensity is the average attention over the announcement day and the two days after, minus its average over the week before:

$$n_e = \frac{1}{3}\sum_{d=0}^{2} a_{i,d_e+d} \;-\; \frac{1}{5}\sum_{d=-7}^{-3} a_{i,d_e+d}. \tag{1}$$

The post window covers announcements made late in the evening, which reach the search data on the following day. The pre window ends three days before the announcement so that late-evening announcements and leaks do not enter the baseline. The exact window of Gavriilidis et al. (2026), the two days before the event, is a robustness choice. The

[6]The bundled birth is the Berlin twins of 2019, announced on September 2, four days before Chancellor Merkel's visit to China. The bundled deaths are Ling Ling (Tokyo, 2008, a week before Hu Jintao's state visit to Japan), Kou Kou (Kobe, 2010, two days after the Senkaku collision) and Le Le (Memphis, 2023, announced during the balloon crisis). All four are excluded from the baseline shock series.

announcement date comes from dated sources, never from the location of the peak in the search data.[7] The soft-power shock sums the intensities of the births announced in country i in month t,

$$z_{i,t} = \sum_{e \in \mathcal{B}_{i,t}} n_e, \tag{2}$$

and equals zero in months without a birth. The headline version uses national attention and the national standard deviation. A second version uses the regional series of the zoo where it is usable, and a third counts births without their dose. The baseline excludes the bundled births, births whose search window is too sparse to measure a dose, and later media days of the same cub, such as the first veterinary examination or the release of the first photograph.

Figure 1 shows what a birth does to attention, as the cumulative sum of daily attention relative to the event's own average in the week before, for the 33 births since 2011. Panel A plots attention to the panda. The cumulative path is flat before the announcement, jumps by about 9.2 standard deviations on the day of the announcement, and then rises at about half a standard deviation a day for three months as the cub stays in the news with its first examination, its sex and its name. Nothing moves before the day of the announcement.[8] Deaths produce a somewhat smaller spike that disappears within a week, and namings and cub debuts a much smaller one, since attention is already elevated after the birth (Appendix Figure A.4, which plots the daily means). Panel B plots attention to China around the same births. It is flat. Births leave searches for China unchanged on the day of the announcement and over the following three months.

Figure 2 shows the monthly SPS, summed across hosts and labeled by host. The 33 births since 2011 that enter the baseline are spread over 14 countries and fifteen years, with

[7] The distinction matters. Everland announced Fu Bao's birth on July 22, 2020, and Google interest in South Korea jumped that day. The largest value in the window came two and a half weeks later and corresponds to no release we could find. Taking the peak as the announcement would have measured a different and unexplained event.

[8] The average rises on the day before the announcement while the median does not. The difference comes from a few announcements made late in the evening, which fall on the previous day in Google's time zone (UTC).

the largest numbers in the United States and Japan. Their doses are heavy-tailed. A few births in countries with a low and stable baseline, such as Belgium in 2016 and France in 2017, raise attention by several times the median dose, and the first stage does not depend on them (Appendix Table A.4).

3.2 Validity of the soft-power shock

Three assumptions underlie the use of births as soft-power shocks. Births are not predictable from past economic and political conditions. They are uncorrelated with the other shocks that move bilateral outcomes, and they are not bundled with state action. We report the checks of each assumption in turn.

Following Ramey (2016), we check the soft-power shock for serial correlation, forecastability and correlation with other shocks. Table 2 reports the joint tests and Appendix Table A.3 the block-by-block results. The shock is serially uncorrelated within countries. Lagged attention to the panda and to China, host industrial production, the exchange rate against the renminbi, China's leading indicator, geopolitical risk and policy uncertainty do not predict it, block by block or jointly. We include calendar-month effects in these regressions, because births are seasonal by biology and any seasonal series would otherwise appear to predict them, and we compute permutation p-values, because the shock is zero in most country-months and the asymptotic test over-rejects with many regressors. The joint test on all macroeconomic and political series has a permutation p-value of 0.41. Births are seasonal, with 73 percent of the baseline births since 2011 in July to September, the same share as all births in the database (Appendix Table A.3).

Appendix Table A.2 correlates the soft-power shock, defined formally in Section 3.1, with the geopolitical risk index of Caldara and Iacoviello (2022) for each country, the economic policy uncertainty index of Baker et al. (2016), the exchange rate against the renminbi, the oil supply news shocks of Känzig (2021), the carbon policy surprises of Känzig (2026), changes in the VIX and changes in China's composite leading indicator.

Over 2011–2025 every pooled correlation is below 0.03 in absolute value, the correlations of the cross-country sum of the shocks with the common series are below 0.1, and none is significant at conventional levels. Over 2004–2025, the soft-power shock is weakly correlated with the level of geopolitical risk, but the correlation disappears within countries. Hosts with more births have higher average geopolitical risk, and births do not occur at times of elevated risk. The correlations support the assumption that births are unrelated to the geopolitical and policy uncertainty that a political shock would carry.

Births are also not bundled with state action. Bundled births are excluded from the baseline, and the hard political events in our database, from meetings with the Dalai Lama to sanctions and summits, do not move attention to the panda. Their first-stage F-statistic for panda attention is 0.3 (Table 2). Nor are births news about China. Searches for China do not respond to births (Figure 1, Panel B), and the soft-power shock has an F-statistic of 0.1 for attention to China, while hard political events do move it. The design holds the policy fixed. Every birth occurs at a zoo whose pandas are already on loan, so the shock measures added exposure under a given policy and does not identify the effect of the loan decision itself.

The breeding attempts that zoos announce offer one more check. Conditional on an announced attempt, whether it ends in a live birth or in a pseudo-pregnancy, resorption or stillbirth is largely a matter of biology. Neither attention to the panda in the three months before the attempt nor the attention that the announcement of the pregnancy itself received differs between the 22 attempts that produced a cub and the 30 that did not (Appendix Table A.3, Panel D). Last, animals that drew national attention for reasons unrelated to China, such as Harambe in the United States or Marius the giraffe in Denmark, raise searches for themselves and leave searches for China unchanged, measured in the same units as the births (Appendix Figure A.7).

3.3 Econometric framework

We estimate the dynamic effects of attention to the panda on bilateral outcomes with local projections (Jordà 2005) on a panel of host countries, instrumenting attention with the soft-power shock (Stock and Watson 2018). For outcome y of country i and horizon h,

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \lambda_t^h + \mu_{i,m(t)}^h + \beta^h A_{i,t} + \Gamma^h X_{i,t-1} + e_{i,t+h}, \tag{3}$$

where $A_{i,t}$ is the monthly panda index of country i, instrumented by $z_{i,t}$. The terms α_i^h, λ_t^h and $\mu_{i,m(t)}^h$ are country, month and country-by-calendar-month effects, and $X_{i,t-1}$ holds the controls listed in Section 3.4. The country-by-calendar-month effects $\mu_{i,m(t)}^h$ are included for outcomes with strong host-specific seasonality, Chinese visitors and air passengers. The trade and attention specifications use country and month effects, and each figure states its set. The coefficient β^h is the response at horizon h to a shift in attention to the panda that is driven by births. We present the reduced form, the local projection of y on $z_{i,t}$ itself, as the main result, and use the instrumental variable estimates to express the responses per unit of attention.

Three conditions identify β^h. The instrument must be correlated with attention to the panda (relevance) and uncorrelated with the other shocks that move bilateral outcomes in the same month (contemporaneous exogeneity). Because the projection does not rely on a model of the dynamics, the instrument must also be uncorrelated with those shocks at leads and lags (Stock and Watson 2018). The biology of births supports the second condition. The forecastability tests above are diagnostics that support the third, since a birth that could be predicted by past shocks would be correlated with them.

Relevance is strong. Table 2 reports the first stage, a regression of the panda index on the instrument with six lags of the index and the same fixed effects, with standard errors clustered by month. For the headline instrument over 2011–2025 the F-statistic is 60.5, and 78.4 with country-by-calendar-month effects. With one instrument it equals

the effective F-statistic of Montiel Olea and Pflueger (2013). A birth with a dose of one standard deviation raises the monthly index by about 2.3 points. The regional version and the count of births without their dose give strong first stages as well, and the first stage holds when any one country is dropped and when the largest events are set to zero.[9][10] The same first stage for attention to China is near zero, which is why the endogenous variable in equation (3) is attention to the panda. The births shift attention to the symbol, and whether this reaches sentiment toward China is what the responses of bilateral flows reveal.

3.4 Specification and sample

The panel covers the 24 host jurisdictions at a monthly frequency. The instrumental-variable sample runs from January 2011 to December 2025, 180 months and 4,320 country-months. Two scope conditions define it, and we state them here once. First, the sample starts in 2011 because Google's daily data for most births before then are too sparse to measure a dose. The announcement rarely stands out in the daily series, and many windows consist mostly of zero days. The 2004–2025 sample, which uses the earlier births with a measurable dose, is a robustness version (Appendix Table A.3). Second, births at the Wakayama zoo before 2017 are treated as missing doses rather than as small ones. Google's topic entity captures Japanese-script searches poorly before 2017, and Wakayama is a small regional market, so these births leave almost no trace in the national series. We keep them as event dummies and set their dose to zero in the soft-power shock. Births at Ueno Zoo in Tokyo are visible in the national series and stay in the shock series.

The left-hand side of equation (3) is the change in the outcome from the month before the announcement, so that β^h is a cumulative response, at horizons from zero to 24 months.

[9] Appendix Table A.4 reports the first stage of the SPS leaving out one host at a time, setting the largest one, three, five and ten event-months to zero, capping the dose at its 90th percentile and taking its logarithm. The F-statistic stays above conventional thresholds in every case.

[10] The strength of the first stage is partly mechanical, because the intensity of each birth and the monthly index are built from the same Google series, much as the event series and the index of Gavriilidis et al. (2026) count the same newspaper articles. The F-statistic therefore speaks to how precisely the dose is measured, not to the exclusion restriction.

Controls are six lags of the outcome's change and of the SPS, together with lags of the exchange rate against the renminbi, the country's geopolitical risk and policy uncertainty indices and the change in China's leading indicator. They also include the host's total exports and the outcome's growth over the previous 12 months. Standard errors are clustered by month, which allows for correlation across countries in the same month.

4 From Attention to Bilateral Flows

Where does the attention a birth creates go? This section follows it from the host public's searches to the flows between the host and China. The outcomes, arranged from attention to flows, are attention to the panda itself, attention to China, Chinese visitors and trade. Chinese visitors are decisions of Chinese travelers, while the shock is measured on the host public. Every rung is estimated in reduced form on the same soft-power shock, and every figure in this section plots one line, the headline specification. The section closes by explaining why earlier studies found positive effects.

Responses are scaled to a typical birth, one whose dose equals the median of the 33 usable births since 2011, 3.08 national standard deviations of announcement-window attention. Outcomes are in percent, 100 times the log change from the month before the announcement, unless stated otherwise. The lines in the figures are smooth local projections (Barnichon and Brownlees 2019), with 68 and 90 percent bands from standard errors clustered by month. We state every zero as a minimum detectable effect (MDE), the effect that the design would detect with 80 percent power in a two-sided test at the 5 percent level, 2.8 times the standard error.[11] Table 3 collects the rungs, from attention to the panda and to China to Chinese visitors and trade. For each outcome it reports the estimate, its 95 percent confidence interval and the MDE at six and twelve months, the

[11] Lewis and Rao (2015) show that even large advertising experiments leave wide intervals for returns, so that a null is stated as a bound. Abadie (2020) discusses when statistically insignificant estimates are informative.

largest absolute t-statistic of the responses before the announcement, and the numbers of births and observations. The figures show the full responses.

4.1 The first stage

A birth raises the host public's attention to the panda at once, and the rise fades over a year. Figure 3 plots the response of the monthly Google index of attention to the panda, which runs from 0 to 100. In the month of the announcement a typical birth raises the index by about 7.0 points. The response is still significant after three months and gone after a year. The daily data show the same shock at higher resolution. Attention to the panda jumps to 9.2 standard deviations above its pre-announcement level on the day of the announcement (Figure 1), measured at the zoo's region where Google reports it. In the national series that defines the soft-power shock, the median birth has a dose of 3.1.

The first stage is the first positive control of the design. It shows that a birth delivers a large and well-measured dose of attention to China's symbol in the population whose behavior the later rungs measure, and that the monthly estimator recovers it. Conflicts with China do not raise attention to the panda, as the first-stage diagnostics of Section 3.3 found (Table 2).

4.2 Attention to China

The first place in which affection for the symbol would reach the country is the host public's curiosity about China. We measure it in the same search engine, in the same countries and over the same days as the first stage, for six groups of Google topics. These are China itself, Chinese brands (Huawei, Xiaomi, TikTok and SHEIN), travel to China (Chengdu and the Great Wall), learning Chinese (Mandarin), Chinese culture (Chinese New Year and Chinese cuisine), and, as a comparison, three other Asian countries (Japan, South Korea and Thailand). Each topic is rescaled to the host's monthly index of the same topic

and measured in units of its own standard deviation on non-event days, as for the panda topic. Births cluster in July to September, so a window that starts at a birth runs into the start of the school year and the autumn travel season. For each birth we therefore take the same host on the same dates in the nearest years before and after it with no panda event, and subtract the attention in these windows from the attention around the birth. Table 4 reports the differences for three windows fixed in advance, days 0–2, 7–30 and 31–90. The China topic is the primary outcome, the four Chinese topics are secondary outcomes, and the other Asian countries are a placebo. Appendix Figures A.5 and A.6 plot the daily series before and after the adjustment.

No Chinese topic moves. Over the first three days, when attention to the panda is at its peak, the China topic changes by 0.05 standard deviations, with a 95 percent confidence interval from −0.14 to 0.24, and the design would detect a move of a fifth of a standard deviation (Table 3). Over the following month and the following quarter the China topic stays within about 0.1 of zero. None of the four secondary topics differs from zero in any window. This holds after a Holm correction over the twelve secondary tests and under randomization inference that relabels each birth window among its calendar windows (Appendix Table A.8). The daily series of the secondary topics do drift up over the month after a birth (Appendix Figure A.5), but the same dates in years without a birth drift up by the same amount (Appendix Table A.8). These are calendar effects, not responses. The placebo countries, which a birth gives no reason to look up, show a small positive difference because their calendar windows drift down. It measures the noise of the adjustment.[12]

The rung locates the point at which the soft-power shock stops. The attention a birth creates attaches to the panda and does not pass to the country that lent it. The host public learns that a cub was born and does not, in the same week or over the following months, look for Chinese products, destinations, language or news. Warmer feelings toward China

[12]The placebo animals of Section 3.2, five events with a usable dose, are too few for the monthly outcomes that follow. In the daily data they raise searches for themselves and do not raise searches for China (Figure A.7).

need not show up as searches, but we observe no corresponding movement in attention to China, and, as the next subsections show, none in the flows that follow.

4.3 Chinese visitors

Chinese visitors are the flow most often credited to pandas, and a flow in which a change of heart on either side, among the host's public or among Chinese travelers who follow news of their country's pandas abroad, would show within months. We measure monthly arrivals of Chinese residents in the 15 hosts with at least ten years of data, from national tourism statistics and Eurostat, with nights spent where arrivals are not reported. The sample drops the months from February 2020 to January 2023, when China closed outbound travel. Because visitor flows have strong and host-specific seasonal patterns, the specification adds country-by-calendar-month effects, and it includes the outcome's growth over the previous 12 months and regressors for each break in a host's statistical source.

A birth produces no detectable change in Chinese arrivals. Figure 4 plots the response, a nearly straight line that rises slowly from zero. Six months after a typical birth the estimate is −0.2 percent, with a 95 percent confidence interval from −6.3 to 5.9 percent (Table 3).[13]

A year after the birth the estimate is 1.3 percent, with an interval from −4.2 to 6.8 percent.[14] In the same equation, conflicts with China lower the same arrivals sharply, as Section 5 shows.

Air passengers between the host and mainland China, a supply-side counterpart

[13]Host-specific policy regimes toward China could move arrivals in the months after a birth. We collect 139 dated regimes with their sources. They cover the host's visa policy toward Chinese visitors, China's visa-free entry for the host's citizens, free-trade agreements, Chinese group-travel bans and direct flights. Entered as the change $D_{i,t+h} - D_{i,t-1}$ over the response window, they move the visitor path by less than a fifth of a standard error at six and twelve months and narrow its bands by about 14 percent (Appendix Figure A.9). Since regimes adopted after a birth could themselves respond to it, this is a sensitivity check rather than a cleaner estimate. The regimes themselves move visitors strongly. A group-travel ban lowers arrivals by 70 to 80 percent, and China's visa-free entry raises them by 10 to 14 percent.

[14]Three births since 2011 were followed within six months by a conflict or a major summit at the same host. Ueno in 2012 was followed by the Senkaku nationalization, Washington in 2013 by a presidential meeting with the Dalai Lama, and the 2017 birth in France by President Macron's state visit to China. Setting them to zero, or giving every hard event its own window, moves the response at six months to about 2 percent, still inside the band, and leaves trade unchanged (Figure A.8).

available monthly for 12 hosts from Eurostat, the US Department of Transportation and the Australian, British and Taiwanese aviation statistics, do not respond either. The response is −0.3 percent at six months and −0.2 percent at a year, with MDEs of about 2 percent, while a conflict lowers passengers by about 6 percent in the same equation (Appendix Figure A.11). The only movement, a rise at two years, comes from one Belgian birth that coincided with the reopening of Brussels airport after the 2016 attacks.

4.4 Bilateral trade and its composition

Trade is the outcome on which the largest published effects of panda births rest, and it is the rung on which the design is most precise. Figure 5 plots the responses of the host's exports to China, its imports from China, its imports of Chinese consumer goods, and the host's total exports, which serve as a negative control. Each comes from equation (3) in reduced form with country and month fixed effects and the control block of Section 3.4, in the 24 host jurisdictions from 2011 to 2025.

Births do not move bilateral trade. Six months after a typical birth, exports to China and imports from China are within half a percent of zero, and after a year they still are. The estimates are precise. Across the two flows and the two horizons, the MDE lies between 1.5 and 2.3 percent of the flow (Table 3), so an effect of the size estimated from annual gravity regressions, 5 to 7 percent, would be detected easily.[15] The conventional intervals rest on month clustering with 33 identifying births. Reassigning each birth at random within its host's tenure of a breeding pair and within the birth season widens the bands to about ±2 to 3 percent for trade and ±7 to 9 percent for visitors, with permutation p-values between 0.4 and 0.9 (Appendix Table A.7), so the zeros do not depend on the asymptotic

[15]The zero holds across the dose-weighted instrument versions (national and regional dose, the dose capped at its 90th percentile), the fixed-effect sets (country and month, country by calendar month, country by year), the control sets, the samples 2011–2019 and 2017–2025, leaving out any one host or any one birth, and in LP-IV with Anderson–Rubin intervals. Appendix Table A.5 reports them. The unweighted count of births, which is also the instrument for the 2004–2025 sample, gives a different import path at one year. The next paragraph examines it. Adding the host-specific policy regimes toward China of footnote 13 leaves the trade paths unchanged (Appendix Figure A.9).

approximation.

The count of births without their dose gives a different import path. With the count in place of the dose, imports from China are about 4 percent higher a year after a birth, significant at conventional levels, while exports to China are unchanged (Appendix Table A.5). The count gives each of the 33 births the same weight, whereas the dose places most of its weight on a few births with large rises in attention. The positive import association in the unweighted specification is not stronger for births that draw more attention and becomes small once host-year conditions are controlled. Births below and above the median dose give the same per-birth estimate, the association is already present in the months before the announcement, and country-by-year fixed effects absorb most of it (Appendix Figure A.10 and Appendix Table A.5).

The one movement in Figure 5 is a rise in exports of about 1 percent in the month of the announcement, which is gone after three months. It is the continuation of a rise in exports that begins about a year before the announcement, and there is no change in slope at the announcement itself (Figure A.3). We read it as the timing of births in good years for the host's exports rather than as a response to them, a point Section 4.5 takes up.

An aggregate zero could hide a response in the goods most exposed to consumer sentiment, offset by movements in industrial inputs. The third panel of Figure 5 therefore isolates imports of Chinese consumer goods (food, beverages, cosmetics, apparel and toys, among others), built from monthly data at the level of the two-digit Harmonized System chapter. Intermediate goods are in Appendix Figure A.13, and China's share of the host's exports in Appendix Figure A.12. Imports of Chinese consumer goods, where affection for China would show first, do not respond. The smoothed response is −0.3 percent after six months and 0.2 percent after a year, and the MDEs are 2.5 and 1.8 percent. Two sharper versions of the test give the same zero. China's share of the host's imports within consumer chapters, which removes everything common to all origins, does not move, and neither does a triple difference in the European hosts that compares imports from China with

imports of the same chapter from other origins in the same month. We find no systematic response across consumer-goods categories.

4.5 Comparison with earlier estimates

The results of this section differ from published estimates of panda effects that are positive and from official figures that are large. The differences follow from differences in design, in what is compared with what and in what is measured.

The finding that exports to China rise after a birth (Longhi and Morelli 2025) comes from a gravity model that compares a host's exports in birth years with its exports in other years. The same comparison in our own panel gives a number of the same order. With host and year fixed effects, an annual birth-year indicator puts the 24 hosts' exports to China 8 percent higher in a birth year over 2011–2025, and 6 percent higher over 2004–2025 (Appendix Table A.6). The comparison uses our sample and a log-linear specification rather than the PPML gravity model of Longhi and Morelli (2025). The lead and the lag of the indicator, however, are nearly as large as the birth year itself. The monthly event study in Appendix Figure A.2 shows why. Exports are about 7 to 8 percent above the host's baseline both in the twelve months before the announcement and in the twelve months after. The change from the year before the announcement to the year after is −1.0 percent with a standard error of 1.4, indistinguishable from zero.

In our data the annual indicator therefore measures a level gap. Births fall in periods in which the host's exports to China are already high, and a birth-year indicator attributes the whole period to the cub. The monthly projection measures the change from the month before the announcement and finds none. The same pattern appears in changes. Exports to China and the host's total exports grow in the year before a birth is announced, and the path flattens after the announcement with no change in slope (Appendix Figure A.3).

The finding that countries hosting pandas receive more Chinese tourists, increasingly after three years of hosting (Okafor et al. 2021), compares hosts with countries that have

no pandas. Hosts differ from other destinations in income, distance and diplomatic ties with China, and China's choice of partners reflects such ties (Buckingham et al. 2013). Our host fixed effects absorb every such level difference and identify from the timing of births within a host, which is when a panda is most in the news. The within-host response is the one in Figure 4, and it is zero.

The official figures measure a different outcome, and they are consistent with our zeros. The Edinburgh assessment, the Adelaide forecast and the Japanese estimate count spending by visitors to the zoo and the city, and the Edinburgh assessment states that its estimates do not cover trade.[16] A panda can raise attendance and spending at the zoo without moving anything between the host and China. Every official figure we could verify measures spending at the zoo, and our outcomes measure exchange with China.

5 Hard Political Events in the Same Panel

This section examines how the same bilateral outcomes respond to political conflicts with China, as a comparison with their response to the soft-power shock. It describes the hard political events, shows their effects on trade and Chinese visitors in the same equations as the births, and documents the largest episodes product by product.

5.1 Hard political events

The database records 98 dated hard political events between China and the host jurisdictions and a few other countries. Of these, 54 are conflicts, coded −1. They are official meetings of heads of state or government with the Dalai Lama, of which there are 20, sanctions and boycotts, disputes over Taiwan and territory, arrests, and conflicts over human rights. The other 44 are summits, state visits and agreements, coded +1. Each event is dated by its first

[16]The estimate for the cub at Ueno Zoo in Kansai University (2017) is a gross economic-effect figure. The Adelaide forecast (Australasian Leisure Management 2009) adds visitor spending on accommodation and hospitality, zoo revenue and jobs.

public announcement, and its intensity is measured, like that of a birth, by the change in the host public's attention, here to the China topic, in the days after the announcement.

The events enter the local projections in two ways. In the trade equations we use their intensity signed by direction, so that each event is weighted by the attention it drew. In the visitor equations we use the count of conflicts in the host and month, because the conflicts that hit Chinese travel hardest, such as the Senkaku nationalization, drew little search attention in the host. Both enter the same equations as the soft-power shock, with their own lags.

These events are not exogenous. Governments choose them, often in response to the state of relations and of trade, and their effects on trade have been documented with careful designs (Fuchs and Klann 2013; Davis et al. 2019). We do not claim to identify them. Because they are chosen by governments, they serve as a comparison, estimated with the same specification, sample and estimator as the births.

5.2 Responses to births and conflicts

Conflicts move the outcomes that births do not. Figure 6 plots the two responses from the same equation, the birth in navy and the conflict in red. In Panel A, a conflict of median intensity lowers exports to China by about 0.2 percent three to six months after the event, significant at the 5 percent level, and the effect fades after a year. Imports from China fall by a similar amount in the first months after a conflict. The comparison rests on exports, which show no movement before conflicts. The birth path in the same equation repeats Figure 5, with the early rise that continues the upswing before the announcement and no response after it.

The visitor response is larger. In Panel B, one conflict with China lowers Chinese arrivals in the host by 17.5 percent after six months, significant at the 5 percent level, and arrivals are still 15.7 percent lower after a year before recovering in the second year. A typical birth in the same equation changes arrivals by −1.1 percent after six months. The

figure compares the two shocks in the same panel, with the same fixed effects and controls. A large and clean shock to attention to China's symbol produces no detectable change in Chinese visitors, while a conflict with the Chinese state reduces them by a sixth.

5.3 Sanction episodes by product

The panel coefficients average over many conflicts that were mostly diplomatic. China's trade sanctions are few, and they are targeted, so their effects are clearest product by product. Table 5 reports, for each major episode, the change in the host's exports to China in the targeted chapter over the following year, relative to the median change in the other hosts, and the host's rank among all hosts. After China imposed anti-dumping deposits on Australian wine in November 2020, Australia's beverage exports to China fell by 290 log points relative to other hosts, the largest fall of the 23 hosts with data. Taiwan's fish exports fell by 326 log points after the Pelosi visit of 2022, and Japan's fish exports by 455 log points after the seafood ban of 2023. Canada's meat exports fell after the arrest of Meng Wanzhou. In most episodes total exports move much less, because the sanctions were aimed at chosen chapters.

Consumer boycotts work through different products. After the Senkaku nationalization of 2012, Japan's vehicle exports to China and Chinese visitors to Japan fell by more than those of any other host, and South Korea's vehicle exports and Chinese visitors fell more than in almost any other host after the THAAD decision and the 2017 boycott and tour ban. South Korean cosmetics, the category most associated with Chinese consumers' taste for Korea, fell only slightly relative to other hosts during the boycott. The larger responses in tourism and in the targeted products are consistent with the targeted nature of the restrictions. Summits, coded +1, do not raise visitors or trade in the panel (Figures A.14 and A.15).

The pattern is consistent with the second channel of the introduction. The flows that fall after a conflict are those the Chinese state controls or can direct. These are customs

clearance of named products, approval of tour groups and the purchases of state firms. A shift in public sentiment in the host has no comparable lever, and the attention a birth creates does not become one.

6 Conclusion

This paper asks whether the attention that cultural diplomacy creates in a foreign public produces economic exchange without an accompanying change in policy. We study the births of giant pandas on loan to foreign zoos, whose timing is set by biology, and measure the dose of each birth in the host public's Google searches. A birth raises attention to the panda. We find no corresponding movement in attention to China, in Chinese visitors or in trade, and the trade estimates are precise enough to exclude effects of the size reported in earlier work.

The evidence separates attention to a country's symbol from economic influence. The symbol draws a foreign public's attention, and the flows between the host and China do not follow it. Conflicts with the Chinese state move the same flows in the same panel, which is consistent with exchange responding to the decisions of governments. The design does not estimate the effect of the loan decision, which comes bundled with state visits and agreements. Nor does it estimate local effects, such as spending at the zoo and in its city, which the national flows we study do not measure.

Acknowledgements

The soft-power shock series and the event database are posted at https://eric-zw-yeh.github.io/panda-shocks/.

Funding. This work was supported by the National Science and Technology Council, Taiwan [grant number 115-2410-H-194-121].

Declaration of competing interest. The author declares no competing interests.

Data availability. The soft-power shock series and the event database are posted at https://eric-zw-yeh.github.io/panda-shocks/. The outcome and control series are from the public sources named in the text.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used Claude (Anthropic) to check the analysis code and to proofread the text. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

Abadie, A. (2020). Statistical nonsignificance in empirical economics. *American Economic Review: Insights 2*(2), 193–208.

Australasian Leisure Management (2009). Adelaide pandas bring $600 million boost. December 2, 2009, https://www.ausleisure.com.au/news/adelaide-pandas-bring-600-million-boost. Accessed October 31, 2026.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636.

Barnichon, R. and C. Brownlees (2019). Impulse response estimation by smooth local projections. *Review of Economics and Statistics 101*(3), 522–530.

Buckingham, K. C., J. N. W. David, and P. Jepson (2013). Diplomats and refugees: Panda diplomacy, soft "cuddly" power, and the new trajectory in panda conservation. *Environmental Practice 15*(3), 262–270.

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225.

China Daily (2023). Woguo lü wai daxiongmao zongti jiankang zhuangkuang lianghao, daxiongmao guoji hezuo chengxiao xianzhu [giant pandas living abroad are in good health; international panda cooperation has achieved notable results]. October 23, 2023, https://cn.chinadaily.com.cn/a/202310/23/WS6536326ba310d5acd876b5c1.html. Accessed October 31, 2026.

Da, Z., J. Engelberg, and P. Gao (2011). In search of attention. *Journal of Finance 66*(5), 1461–1499.

Davis, C. L., A. Fuchs, and K. Johnson (2019). State control and the effects of foreign relations on bilateral trade. *Journal of Conflict Resolution 63*(2), 405–438.

Disdier, A.-C. and T. Mayer (2007). Je t'aime, moi non plus: Bilateral opinions and international trade. *European Journal of Political Economy 23*(4), 1140–1159.

Du, Y., J. Ju, C. D. Ramírez, and X. Yao (2017). Bilateral trade and shocks in political relations: Evidence from China and some of its major trading partners, 1990–2013. *Journal of International Economics 108*, 211–225.

Fisman, R., Y. Hamao, and Y. Wang (2014). Nationalism and economic exchange: Evidence from shocks to Sino-Japanese relations. *Review of Financial Studies 27*(9), 2626–2660.

Fuchs, A. and N.-H. Klann (2013). Paying a visit: The Dalai Lama effect on international trade. *Journal of International Economics 91*(1), 164–177.

Gavriilidis, K., D. R. Känzig, R. Raghavan, and J. H. Stock (2026). The macroeconomic effects of climate policy uncertainty. NBER Working Paper 34762, National Bureau of Economic Research.

Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange? *Quarterly Journal of Economics 124*(3), 1095–1131.

Hartig, F. (2013). Panda diplomacy: The cutest part of China's public diplomacy. *The Hague Journal of Diplomacy 8*(1), 49–78.

Heilmann, K. (2016). Does political conflict hurt trade? Evidence from consumer boycotts. *Journal of International Economics 99*, 179–191.

Hong Kong Government News (2024). Da xiongmao an an ke ke di gang [giant pandas an an and ke ke arrive in Hong Kong]. September 26, 2024, https://www.news.gov.hk/chi/2024/09/20240926/20240926_120141_076.html. Accessed November 1, 2026.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Kansai University (2017). Miyamoto katsuhiro meiyo kyoju ga suitei: Ueno dobutsuen no panda no akachan tanjo no keizai koka [professor emeritus katsuhiro miyamoto estimates the economic effect of the birth of a panda cub at ueno zoo]. Press release No. 12, June 7, 2017, https://www.kansai-u.ac.jp/global/guide/pressrelease/2017/No12.pdf. Forecast issued before the birth: about 26.7 billion yen a year in Tokyo if the cub were born. Accessed November 1, 2026.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Känzig, D. R. (2026). The unequal economic consequences of carbon pricing. *American Economic Review 116*(11), 4003–4040.

Lagerborg, A., E. Pappa, and M. O. Ravn (2023). Sentimental business cycles. *Review of Economic Studies 90*(3), 1358–1393.

Lewis, R. A. and J. M. Rao (2015). The unfavorable economics of measuring the returns to advertising. *Quarterly Journal of Economics 130*(4), 1941–1973.

Liu, C. (2024). New pair of giant pandas to return to Washington DC. *Global Times*, May 30, 2024, https://www.globaltimes.cn/page/202405/1313357.shtml. Accessed October 31, 2026.

Longhi, M. and C. Morelli (2025). The role of Chinese panda ambassadors in international trade. DEMS Working Paper 550, University of Milano-Bicocca, Department of Economics, Management and Statistics. Revised July 2026.

Montiel Olea, J. L. and C. Pflueger (2013). A robust test for weak instruments. *Journal of Business and Economic Statistics 31*(3), 358–369.

NBC 7 San Diego (2024). The nearly 30-year history between giant pandas and the San Diego Zoo. August 8, 2024, https://www.nbcsandiego.com/news/local/timeline-giant-pandas-at-san-diego-zoo/3503904/. Accessed November 1, 2026.

Nye, Jr., J. S. (2004). *Soft Power: The Means to Success in World Politics*. New York: PublicAffairs.

Office of the Governor of California (2024). Giant pandas set to return to California. February 23, 2024, https://www.gov.ca.gov/2024/02/23/giant-pandas-set-to-return-to-california/. Accessed November 1, 2026.

Okafor, L. E., C. M. Tan, and U. Khalid (2021). Does panda diplomacy promote Chinese outbound tourism flows? *Journal of Hospitality and Tourism Management 49*, 54–64.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, Volume 2, pp. 71–162. Elsevier.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Rose, A. K. (2016). Like me, buy me: The effect of soft power on exports. *Economics & Politics 28*(2), 216–232.

Scottish Enterprise (2012). Tourism economic impact assessment of the UK giant panda project: Summary report. Report SC-8062, Scottish Enterprise. https://scottish-enterprise.com/media/l1eaejdz/tourism-economic-impact-assessment-of-the-uk-giant-panda-project-summary-report.pdf, accessed October 31, 2026.

Smithsonian's National Zoo and Conservation Biology Institute (2025). The history of giant pandas at the Smithsonian's national zoo and conservation biology institute. https://nationalzoo.si.edu/animals/history-giant-pandas-zoo. Accessed November 1, 2026.

Stock, J. H. and M. W. Watson (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *Economic Journal 128*(610), 917–948.

Willis, E. L., D. C. Kersey, B. S. Durrant, and A. J. Kouba (2011). The acute phase protein ceruloplasmin as a non-invasive marker of pseudopregnancy, pregnancy, and pregnancy loss in the giant panda. *PLoS ONE 6*(7), e21159.

Xi, J. (2023). Full text of Xi's speech at welcome dinner by friendly organizations in the United States. People's Daily Online (Xinhua), November 16, 2023, https://en.people.cn/n3/2023/1116/c90000-20098529.html. Speech delivered in San Francisco on November 15, 2023; accessed November 1, 2026.

Zhang, H., D. Li, C. Wang, and V. Hull (2009). Delayed implantation in giant pandas: The first comprehensive empirical evidence. *Reproduction 138*(6), 979–986.

Table 1: The Panda Event Database: Events by Type

This table classifies the 491 events in the database by type. An event is dated by its first public announcement (a zoo release, a wire story or a national broadcaster); "Day" is the share of events whose announcement is known to the day. "2004+" counts events announced from February 2004, when Google's daily data begin. "Bundled" counts events coded +1 (a state visit, summit or major agreement between the host and China within seven days of the announcement) and −1 (a diplomatic conflict, sanction or recall in the same window). "Primary" is the share of events with at least one primary source. "Pregnancy" includes announcements of artificial insemination and mating. Dates run from August 1980 (Mexico, gift era) to January 2026. Appendix Table A.1 lists the events by host.

Type	Events	Day	2004+	Bundled +1	Bundled −1	Primary
Loan announcement	25	0.84	25	18	0	0.92
Agreement	17	0.82	17	6	0	1.00
Arrival	45	0.93	30	4	0	0.96
Public debut	30	0.90	26	4	0	0.80
Pregnancy	75	0.95	75	0	0	0.97
Birth	67	1.00	56	1	0	0.94
Failed pregnancy	24	0.96	24	0	0	0.96
Naming	46	0.96	44	3	0	0.87
Cub debut	54	0.91	50	0	0	0.85
Cub departure	33	0.94	33	1	0	0.82
Death	24	0.83	16	1	2	0.79
Return	21	0.71	18	0	2	0.86
Recall	2	1.00	2	0	1	1.00
Extension	28	0.82	28	2	0	0.82
Total	491	0.91	444	40	5	0.90

Table 2: The Soft-Power Shock: First Stage and Forecastability

Panel A reports first-stage regressions of the monthly Google index of attention to the giant panda (0–100) or to China on a shock series, six lags of the dependent variable, country and month fixed effects (CT) and, in the last column, country-by-calendar-month fixed effects (CT+CM); standard errors are clustered by month, and with one instrument the cluster-robust F-statistic equals the effective F-statistic of Montiel Olea and Pflueger (2013). The headline instrument is the first row. The hard political events (China-topic intensity times the direction of the event, +1 rapprochement, −1 conflict) are not used as an instrument; the rows that use them are diagnostics. The coefficient is in index points per standard deviation of dose. Sample 2011–2025, 4,320 country-months, 180 months. Panel B regresses the soft-power shock on six lags of each variable with country and calendar-month fixed effects; F is the Wald statistic for the block, with its asymptotic p-value and a permutation p-value from 300 draws that reshuffle the shock across years within country-by-calendar-month cells. The last row is the AR(1) coefficient of the shock with country fixed effects. Appendix Table A.3 reports the alternative samples and versions of the instrument, the block-by-block forecastability tests, the seasonality of births and the comparison of successful and failed breeding attempts.

Panel A: First stage

Dependent	Shock series	Coef.	s.e.	F (CT)	F (CT+CM)	Event-months
Panda attention	Births, national dose	2.28	0.29	60.51	78.36	33
China attention	Births, national dose	0.03	0.09	0.10	0.59	33
Panda attention	Hard political events	−0.06	0.11	0.28	0.26	69
China attention	Hard political events	−1.32	0.54	5.97	5.53	69

Panel B: Forecastability of the soft-power shock, 2011–2025

Block of lagged variables	Regressors	F	p (asymptotic)	p (permutation)
All macro and political series jointly	24	2.54	0.00	0.41
Attention, FX, CLI and GPR jointly	30	2.91	0.00	0.36
AR(1) coefficient		−0.00	0.63	

Table 3: Responses to One Typical Panda Birth: Estimates, Confidence Intervals and Minimum Detectable Effects

Response of each outcome to the soft-power shock of one typical panda birth ($z^{\text{nat}}_{\text{birth}} = 3.08$ national-SD units of announcement-window attention; all births kept), h months after the announcement. Estimate: the ordinary local-projection coefficient. 95% CI: the conventional interval, the estimate plus and minus 1.96 standard errors, with standard errors clustered by month. MDE: the minimum detectable effect, 2.8 times the standard error, the effect the design detects with 80% power in a two-sided test at the 5% level. Appendix Table A.7 reports permutation p-values and null bands that reassign each birth to a random month of its host's tenure of a breeding pair. The smoothed responses are in Figures 3 to 5. Pre-trend: largest $|t|$ over $h = -12, \ldots, -1$ of $y_{i,t+h} - y_{i,t-13}$ without the control block. Rung 1a: national Google panda-topic index, 24 hosts, country and month fixed effects, macro controls. Rung 1b: daily China-topic index around the 33 births, mean of days 0–2 minus days −7 to −3 in host-SD units, event-clustered standard error. Rung 2, visitors: 15 hosts including Taiwan, country, month and country×calendar-month fixed effects, months 2020-02 to 2023-01 dropped, control block with 12-month pre-growth, source-break regressors; 15 births identify the path. Rung 2, consumer goods: host imports from China in HS2 consumer chapters, 24 hosts, country and month fixed effects, control block. Rung 3: IMF bilateral trade, 24 hosts, country and month fixed effects, control block. Shock months 2011–2025. Obs.: rows at $h = 6$.

		$h = 6$			$h = 12$			Pre-trend	N	
Rung and outcome	Units	Estimate	95% CI	MDE	Estimate	95% CI	MDE	max $\|t\|$	births	obs.
1a. Panda attention (A_{panda})	index pts	1.70	[−0.51, 3.92]	3.17	0.26	[−0.76, 1.29]	1.46	2.15	33	4,176
1b. China attention (daily)	host SD	days 0–2: estimate 0.02, 95% CI [−0.12, 0.17], MDE 0.21						–	33	33
2. Chinese visitors	%	−0.16	[−6.26, 5.94]	8.72	1.30	[−4.22, 6.82]	7.89	1.22	15	1,785
2. Consumer imports from China	%	−0.79	[−2.55, 0.97]	2.51	0.41	[−0.81, 1.64]	1.75	1.20	33	3,768
3. Exports to China	%	−0.28	[−1.89, 1.32]	2.29	0.16	[−0.88, 1.20]	1.48	2.44	33	4,176
3. Imports from China	%	−0.40	[−1.79, 0.98]	1.97	0.33	[−0.89, 1.56]	1.75	2.76	33	4,176

Table 4: Attention to China around Panda Births, Net of the Same Dates in Years without a Panda Event

Each cell averages, across births, the birth-minus-placebo difference in daily Google attention, in host standard deviations, over the days of the window. Per birth and topic, daily attention is rescaled to the host's monthly index, divided by the host's standard deviation and demeaned over the window's own days −37 to −7; the same series for the one or two calendar placebo windows of the birth (same host, same month-day, nearest earlier and later year in 2011–2025 with no panda event of the host within 60 days of day 0) is averaged and subtracted. Differences are averaged over the group's topics usable in both windows, on the unsmoothed daily series; the Giant panda topic is excluded. The three windows are fixed in advance. China (country topic) is the primary outcome, the next four groups are secondary and Other Asian countries (Japan, South Korea, Thailand) is a placebo group. Est.: cross-birth mean. 95% CI and two-sided p: from a t distribution with $n - 1$ degrees of freedom. Births with a matched placebo window: China 33, Chinese brands 32, Travel to China 18, Learning Chinese 17, Chinese culture 31, Other Asian countries 33. Appendix Table A.8 reports multiple-testing and randomization-inference p-values and the birth windows without the placebo subtraction; Appendix Figures A.5 and A.6 plot the daily series.

	Days 0–2			Days 7–30			Days 31–90		
Group	Est.	95% CI	p	Est.	95% CI	p	Est.	95% CI	p
China	0.05	[−0.14, 0.24]	0.60	−0.10	[−0.25, 0.04]	0.17	−0.11	[−0.36, 0.14]	0.36
Chinese brands	0.19	[−0.11, 0.49]	0.20	0.12	[−0.14, 0.39]	0.35	0.13	[−0.39, 0.64]	0.62
Travel to China	0.10	[−0.22, 0.41]	0.53	0.10	[−0.19, 0.40]	0.47	−0.05	[−0.35, 0.25]	0.74
Learning Chinese	−0.24	[−0.70, 0.22]	0.29	−0.24	[−0.70, 0.22]	0.28	−0.06	[−0.59, 0.48]	0.82
Chinese culture	0.03	[−0.25, 0.32]	0.82	−0.02	[−0.19, 0.15]	0.78	0.02	[−0.21, 0.24]	0.89
Other Asian countries	0.31	[0.02, 0.60]	0.04	0.21	[0.00, 0.41]	0.05	0.21	[−0.02, 0.44]	0.07

Table 5: Sanction Episodes: Exports to China in the Targeted Chapter, in Total, and Chinese Visitors

For each episode (month t of the announcement), the host's mean log value over $t+1, \ldots, t+12$ minus its mean over $t-12, \ldots, t-1$, times 100, minus the median of the same change across the other hosts with data. In parentheses: the host's rank among the hosts with data (1 = largest fall) and the number of hosts. Chapter and total exports: host exports to China in US dollars (UN Comtrade; Eurostat Comext for EU members; China's mirror records for Taiwan). Visitors: Chinese visitor arrivals (nights for AT, FI, NL, DK; air entries for MX) in the 15 tourism hosts; – = host not in the tourism panel. [b] A recording break in the host's visitor series falls inside the window. [c] The window touches China's outbound closure (2020-02 to 2023-01). Vehicles (chapter 87) from the capital-goods chapter file; food = chapters 02–04, 07–09, 16–21. Descriptive statistics, not estimates.

				Change over the following year, log points (rank)		
Episode	Host	Date	Targeted chapter	Chapter exports	Total exports	Visitors
Senkaku nationalisation, boycotts	JP	2012-09	87 vehicles	−41 (1/20)	−23 (2/20)	−50 (1/15)
			all consumer goods	−35 (1/20)		
			all capital goods	−23 (1/20)		
THAAD deployment decision	KR	2016-07	87 vehicles	−66 (1/24)	−9 (6/24)	−33 (2/15)
			33 cosmetics	−8 (9/23)		
			food	−8 (6/23)		
Lotte land swap, boycott, tour ban	KR	2017-02	87 vehicles	−77 (1/24)	−7 (6/24)	−90 (1/15)
			food	−31 (1/23)		
			33 cosmetics	−16 (7/23)		
Section 301 round 1, retaliation	US	2018-07	08 fruit	−41 (5/21)	−23 (2/24)	−19 (1/15)
			food	−37 (2/23)		
			02 meat	−36 (5/20)		
Meng Wanzhou arrest, meat/canola	CA	2018-12	02 meat	−173 (3/21)	−14 (3/24)	−7 (5/15)[b]
			food	−13 (4/23)		
Barley tariffs, beef suspensions	AU	2020-05	22 beverages (wine)	−134 (1/23)	−6 (7/24)	–[c]
			02 meat	−80 (4/20)		
			food	−44 (3/23)		
Huawei 5G ban	GB	2020-07	all consumer goods	−9 (11/24)	−31 (2/24)	–
			all capital goods	+3 (16/24)		
Wine anti-dumping deposits	AU	2020-11	22 beverages (wine)	−290 (1/23)	−4 (8/24)	–[c]
Pelosi visit, fruit and fish bans	TW	2022-08	03 fish	−326 (1/20)	−19 (3/23)	–[c]
			08 fruit	−173 (2/20)		
Fukushima water release, seafood ban	JP	2023-08	03 fish	−455 (3/20)	+1 (13/23)	–[c]
			food	−107 (1/22)		

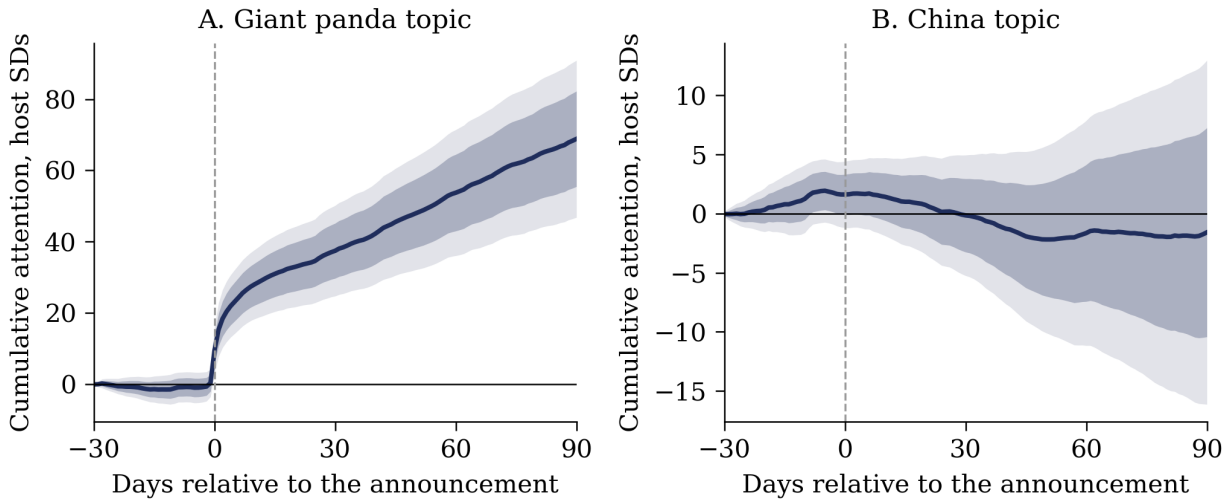


Figure 1: **Attention around Panda Births.** Each panel plots the cumulative sum of daily Google attention in the host country around the 33 baseline births announced from January 2011 to December 2025, in units of the country's standard deviation of daily attention on non-event days, minus the event's own mean over days −7 to −3, cumulated from day −30 and averaged across events; days with no usable value count as zero. Day 0 is the first public announcement. Panel A: giant panda topic, measured in the region of the zoo where Google's regional data are usable and nationally otherwise. Panel B: China topic, national. Dark and light bands are 68 and 90 percent intervals from cross-event standard errors of the cumulative sum. Appendix Figure A.4 plots the daily means, with deaths and namings.

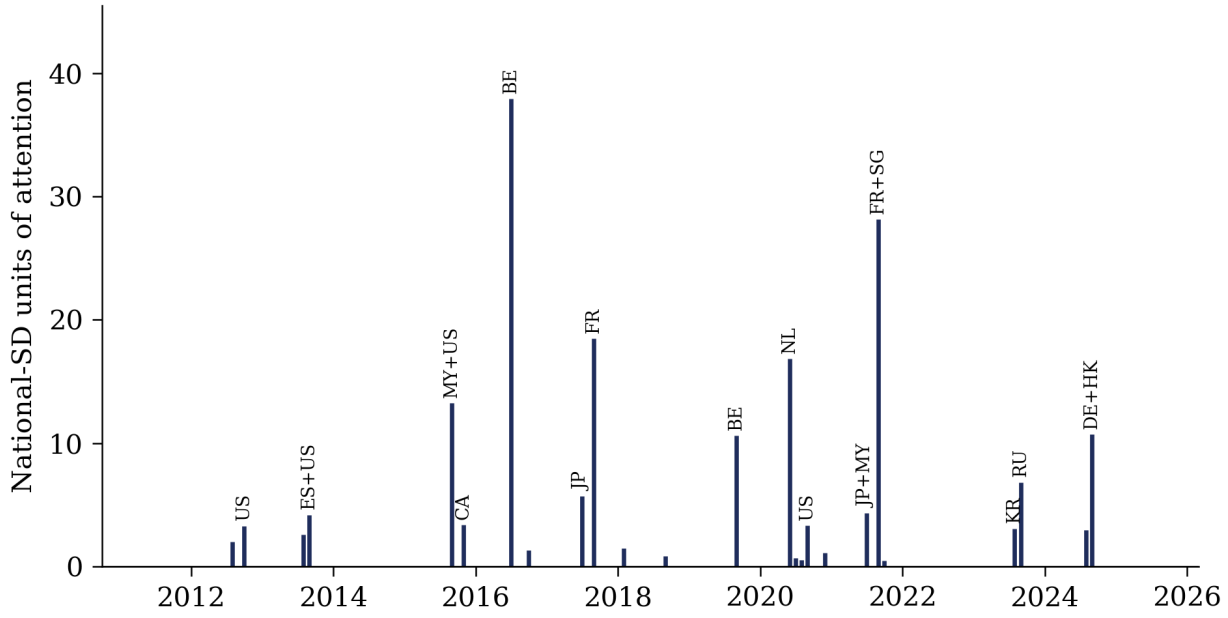


Figure 2: **The Soft-Power Shock.** The figure plots the monthly soft-power shock $z_{i,t}$ of equation (2), the sum of the national-dose intensities of the births announced in host i in month t, summed across hosts, in national standard deviations of announcement-window attention, January 2011 to December 2025. Spikes of at least three units are labeled with the host (two hosts joined by + when their births fall in the same month). The series is zero in months without a usable birth; the 33 nonzero host-months in 14 hosts are the baseline births of Appendix Table A.3, Panel C.

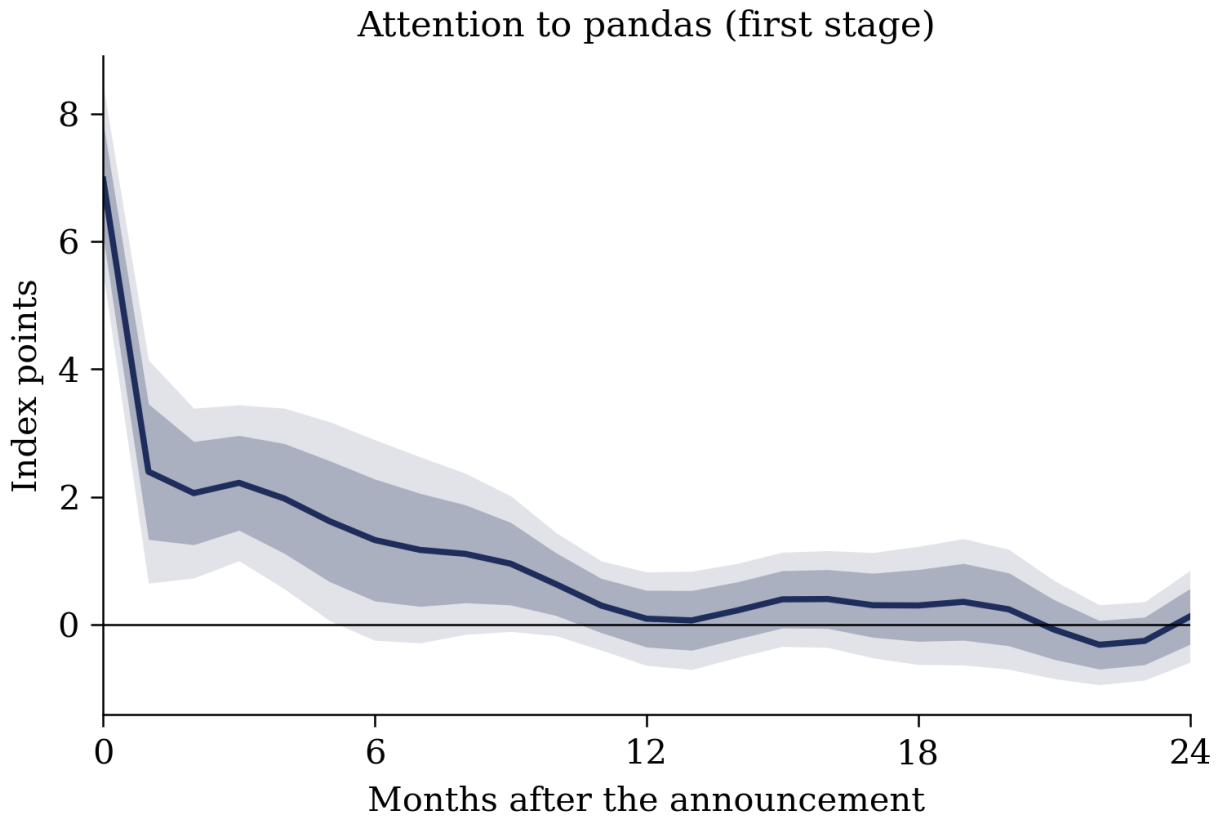


Figure 3: **The First Stage: Attention to the Panda after a Birth.** The line is the response of the monthly Google index of attention to the giant panda in the host (0–100, index points) to one typical birth, a birth with the median dose of the 33 usable births since 2011 (3.08 national standard deviations of announcement-window attention). Local projection of $A_{i,t+h} - A_{i,t-1}$ on the soft-power shock with country and month fixed effects, six lags of the change in the index and of the shock, and lags of the exchange rate against the renminbi, country geopolitical risk and policy uncertainty and China's leading indicator; 24 hosts, shock months 2011–2025. Smooth local projection (Barnichon and Brownlees 2019) with the penalty chosen by cross-validation, which keeps the impact response. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month.

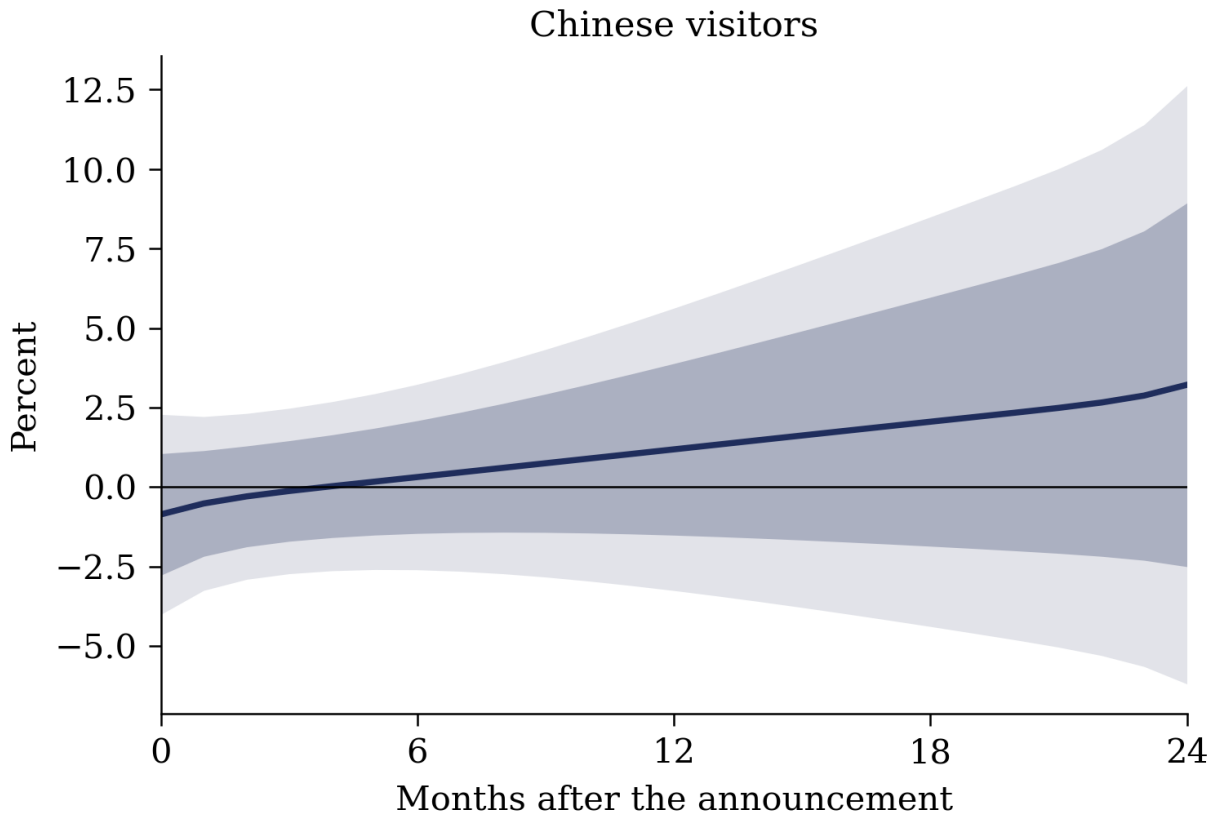


Figure 4: **Chinese Visitors after a Panda Birth.** The line is the response of Chinese visitor arrivals in the host (percent, 100 times the log change from the month before the announcement) to one typical birth. Fifteen hosts with at least 120 months of data: arrivals, nights for Austria, Denmark, Finland and the Netherlands, and air entries of Chinese nationals for Mexico; all births kept, of which 15 identify the path. Country, month and country-by-calendar-month fixed effects; six lags of the outcome's monthly change and of the shock; lags of the exchange rate against the renminbi, country geopolitical risk and policy uncertainty, China's leading indicator and host total exports; the outcome's growth over the previous 12 months; regressors for breaks in the statistical source. Months from February 2020 to January 2023 are dropped. Smooth local projection (Barnichon and Brownlees 2019), penalty chosen by cross-validation. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month; shock months 2011–2025.

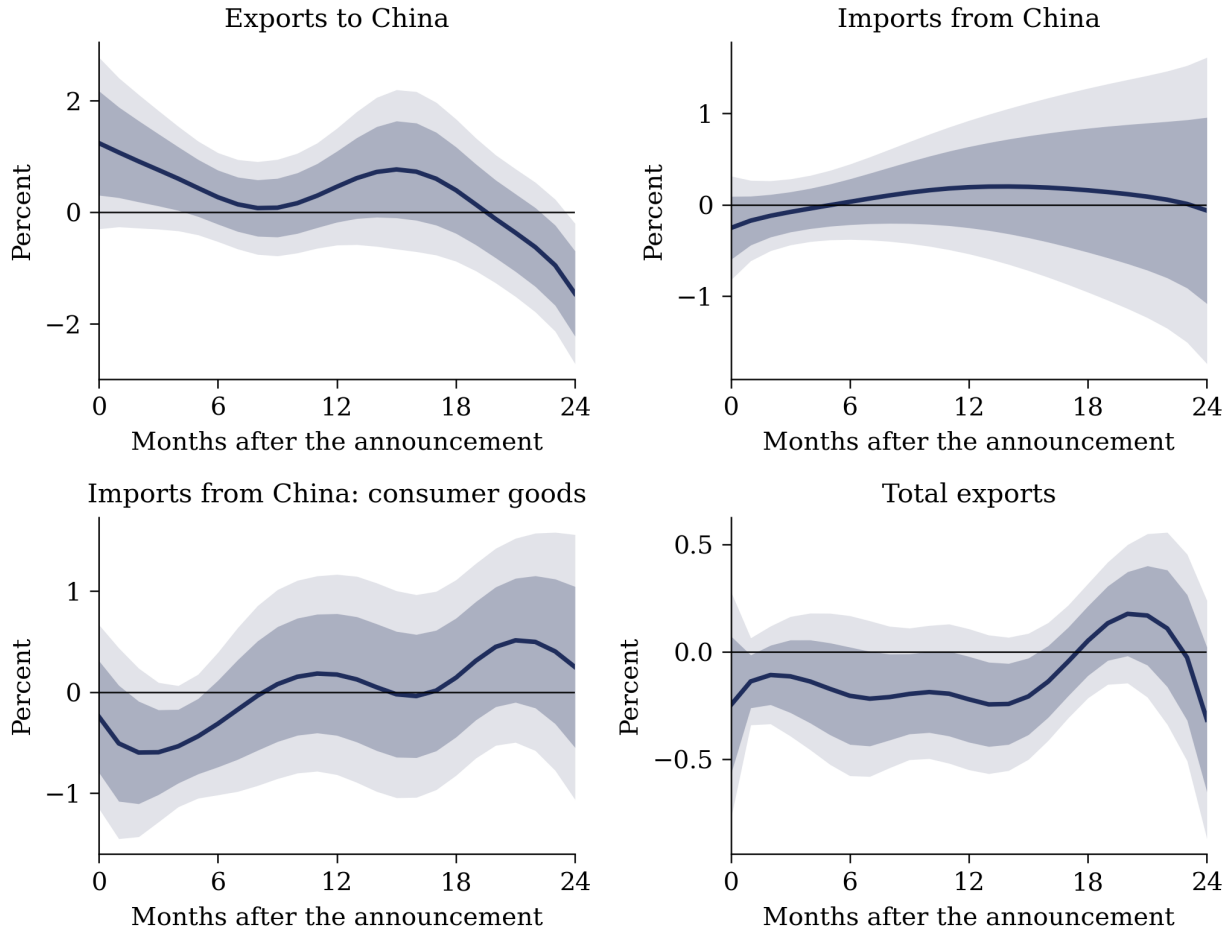
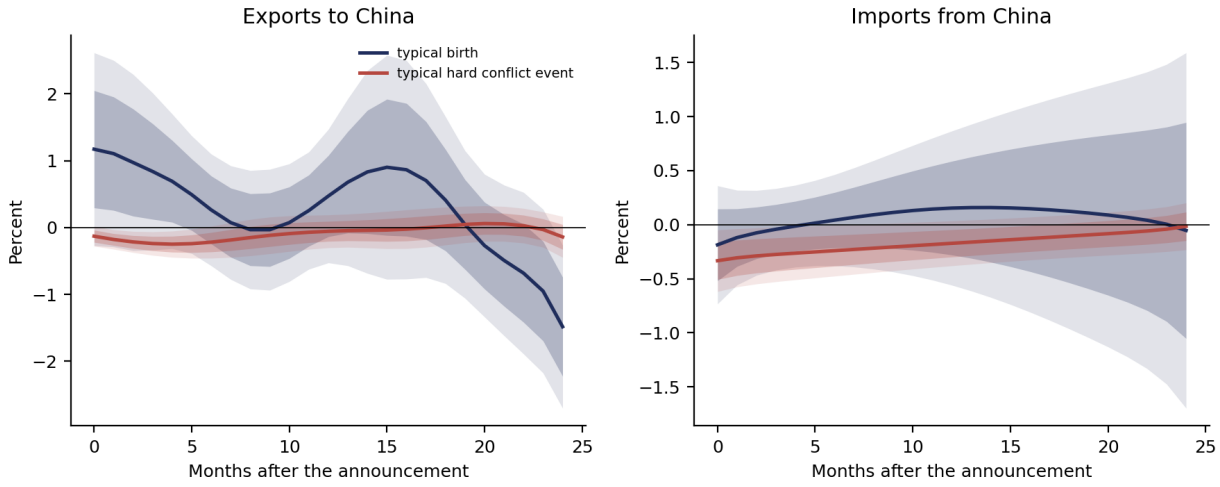


Figure 5: **Trade with China after a Panda Birth.** Each panel plots the response to one typical birth of the host's exports to China, its imports from China, its imports from China in the consumer-goods chapters of the Harmonized System (02–04, 07–09, 16–22, 33, 61–64 and 95), and its total exports (percent, 100 times the log change from the month before the announcement). Country and month fixed effects; six lags of the outcome's monthly change and of the shock; lags of the exchange rate against the renminbi, country geopolitical risk and policy uncertainty, China's leading indicator, host total exports and the other bilateral flow with China; the outcome's growth over the previous 12 months. Aggregate trade is host-reported in US dollars (IMF; Taiwan's Ministry of Finance); consumer goods are from UN Comtrade, Eurostat Comext for the eight EU hosts and China's mirror records for Taiwan. 24 hosts, shock months 2011–2025. Smooth local projections (Barnichon and Brownlees 2019) with the penalty chosen by cross-validation; in the consumer-goods panel the penalty is the larger of the cross-validated penalty and the penalty at which the effective degrees of freedom equal six. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month. China's share of the host's exports and imports of intermediate goods are in Appendix Figures A.12 and A.13.

Panel A: Exports to China and imports from China



Panel B: Chinese visitors

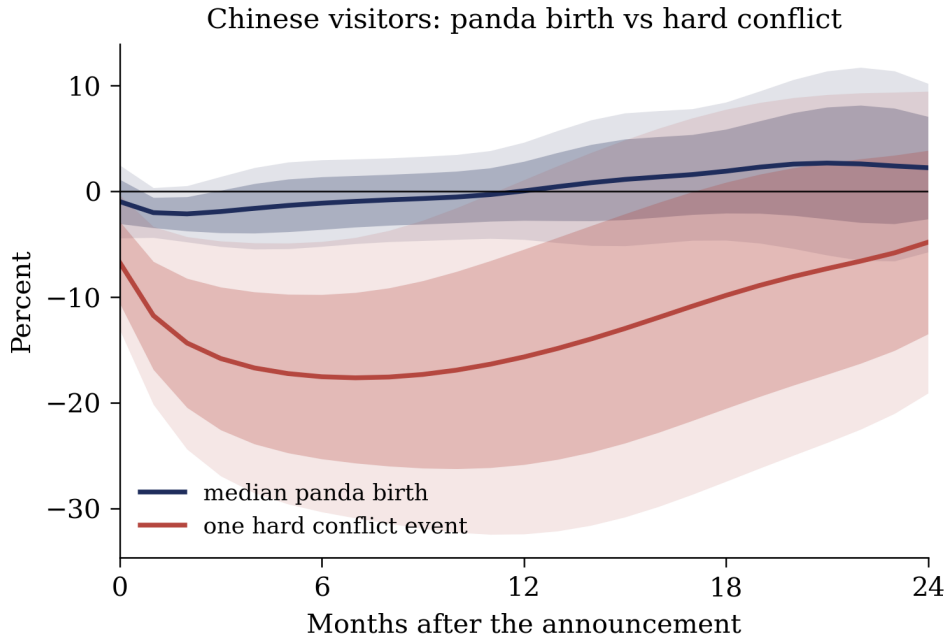


Figure 6: **A Panda Birth and a Conflict with China, in One Equation.** Navy: one typical birth. Red: a conflict with China. Both shocks and six lags of each enter the same local projection. In Panel A the conflict is scaled to the median absolute China-topic intensity of conflict events (signed negative), and a death shock enters the same equation (its path is in Figure A.15); specification as in Figure 5. In Panel B the conflict is one conflict-direction event in the host (a meeting with the Dalai Lama, a sanction or boycott, a dispute over Taiwan or territory, an arrest or a human-rights dispute), entered as a count; specification as in Figure 4, and the two paths are smoothed jointly with the penalty rule of Appendix Figure A.13. Percent. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month.

Appendix A Additional Tables and Figures

This appendix collects the tables and figures referred to in Sections 2 to 5. The exhibits for Section 2 are the attention index of each host and the events by host. The exhibits for Section 3 are the correlations of the soft-power shock with other shocks, its full diagnostics and the first stage leaving out hosts and trimming the dose. The exhibits for Section 4 are the robustness table, the replication bridge to the annual estimate, the randomization inference, the daily attention means, the trade pre-trends, the placebo animals, the attention series net of the calendar and their test by window, the treatment of births followed by hard events, the policy-regime controls, the count of births against the dose, the air-passenger responses and the remaining trade series. The exhibits for Section 5 are the hard-event variations. Each exhibit is one variation of the headline specification, described in its notes.

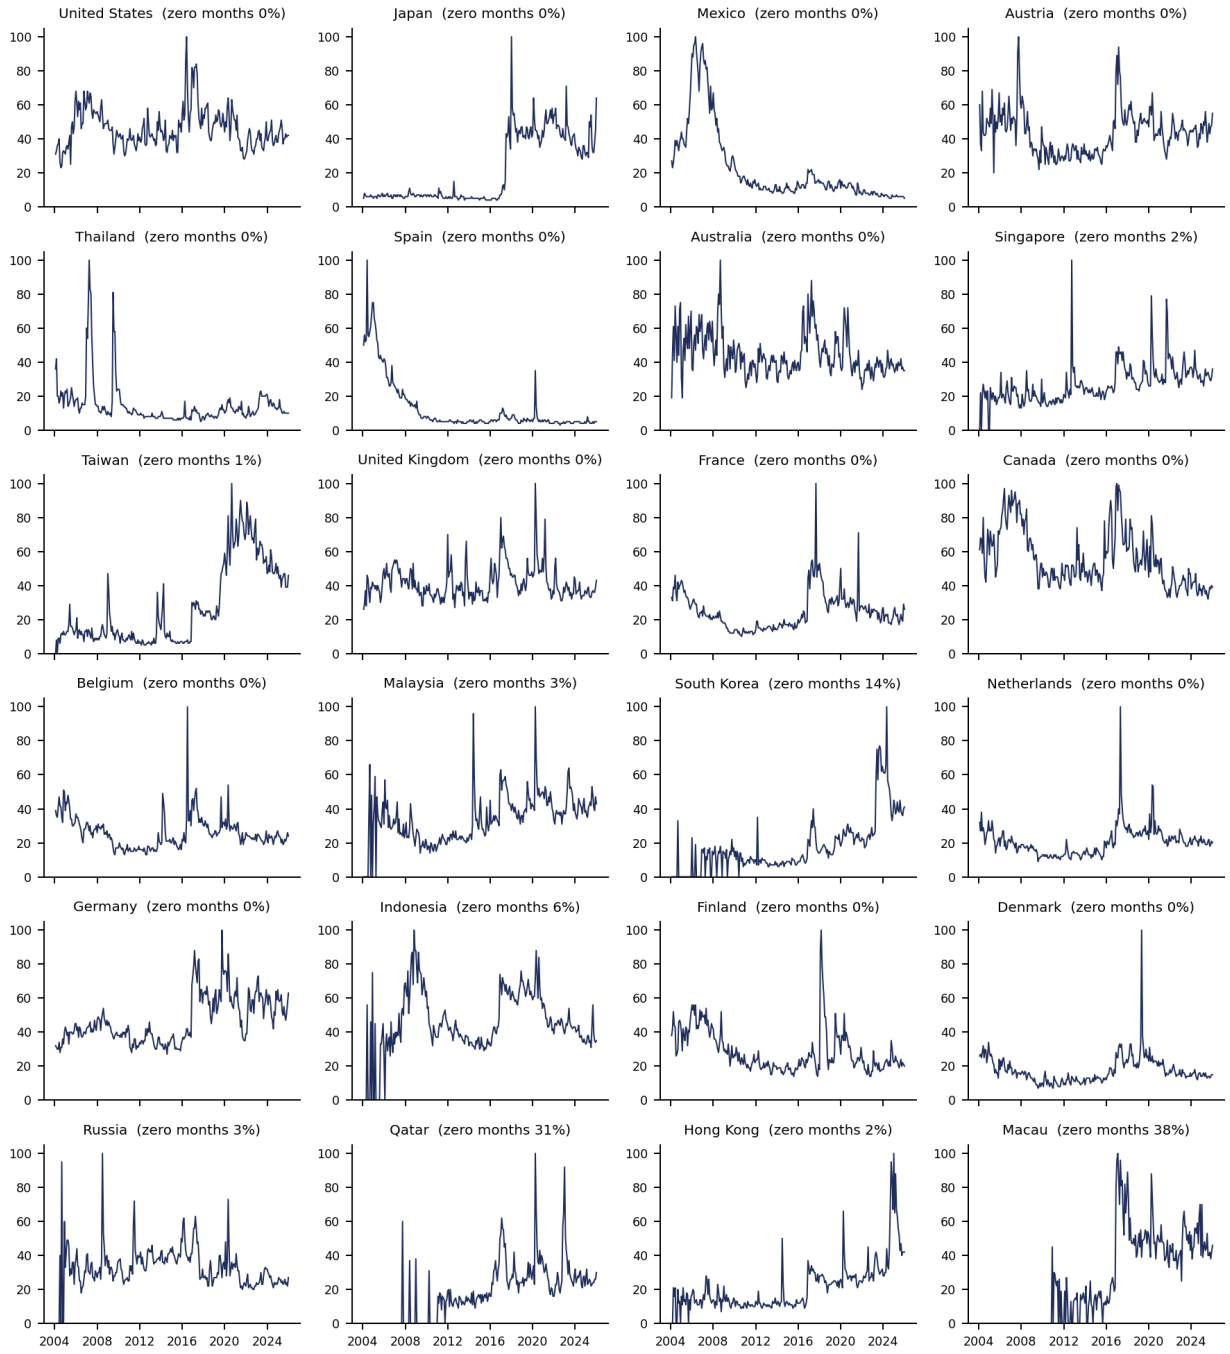


Figure A.1: **Attention to the Giant Panda in the Host Countries, 2004–2025.** Each panel plots the monthly Google Trends index for the giant panda topic in one host jurisdiction, January 2004 to December 2025. Google normalizes each series to 100 at its own maximum, so levels are comparable within a panel but not across panels. Panel titles give the share of months in which Google reports zero interest. Each series is a single pull.

Table A.1: The Panda Event Database: Events by Host

This table counts the events of Table 1 by host jurisdiction, with births and announced breeding attempts (from the separate attempts file); "Day" is the share of events whose announcement is known to the day; "First" and "Last" are the years of the host's first and last event.

Host	Events	Births	Attempts	Day	First	Last
United States	130	16	27	0.92	1996	2025
Japan	78	16	22	0.92	1994	2026
Austria	30	4	5	0.97	2003	2025
Spain	26	4	5	0.81	1996	2024
United Kingdom	19	0	7	1.00	2011	2023
Thailand	18	1	6	0.89	2003	2024
Germany	16	2	3	0.94	1995	2025
Malaysia	16	3	3	0.94	2012	2025
France	15	2	3	1.00	2011	2025
Taiwan	15	2	9	1.00	2005	2022
Canada	14	1	3	1.00	2012	2020
Hong Kong	14	1	5	0.93	1999	2025
South Korea	14	2	3	0.93	1994	2024
Mexico	13	6	1	0.85	1980	2022
Netherlands	13	2	3	0.85	2015	2025
Australia	12	0	7	0.83	2007	2024
Belgium	11	2	3	0.91	2013	2025
Finland	7	0	3	0.86	2016	2024
Macau	7	1	1	1.00	2010	2016
Russia	7	1	1	0.57	2001	2024
Singapore	7	1	7	1.00	2009	2024
Denmark	4	0	6	0.75	2014	2019
Qatar	3	0	0	0.67	2020	2022
Indonesia	2	0	0	0.50	2017	2017
Total	491	67	133	0.91	1980	2026

Table A.2: Correlations of the Soft-Power Shock with Other Indices and Shocks

This table reports Pearson correlations between the soft-power shock $z_{i,t}$ (equation 2, regional dose where usable and national otherwise; the national version has the same event-months) and other series. Panel A uses country-specific series in the country-month panel: the country geopolitical risk index of Caldara and Iacoviello (2022) (GPR), the country economic policy uncertainty index of Baker et al. (2016) (EPU), their monthly changes, and the log change of the host currency against the renminbi. "Pooled" is the correlation across all country-months; "Within" removes country means. Panel B uses series common to all countries: the oil supply news shocks of Känzig (2021), the carbon policy surprises of Känzig (2026) (available to 2019), the monthly change in the VIX and the change in the OECD composite leading indicator for China. "Pooled" is the correlation in the country-month panel; "World sum" is the correlation of $\sum_i z_{i,t}$ with the series. p-values are for the null of zero correlation. GPR is not available for Austria, Macau, Qatar and Singapore; EPU covers 12 hosts.

	2011–2025				2004–2025			
	Pooled	p	Within	p	Pooled	p	Within	p
Panel A: Country-specific series								
GPR	0.02	0.31	−0.01	0.50	0.03	0.04	0.00	0.98
ΔGPR	−0.00	0.77	−0.01	0.76	−0.00	0.94	−0.00	0.93
EPU	−0.02	0.47	−0.02	0.35	−0.00	0.80	−0.01	0.76
ΔEPU	−0.02	0.43	−0.02	0.44	−0.02	0.36	−0.02	0.36
Δ log FX vs. CNY	−0.01	0.52	−0.01	0.55	−0.01	0.44	−0.01	0.47
	Pooled	p	World sum	p	Pooled	p	World sum	p
Panel B: Common series								
Oil supply news shock	−0.00	0.81	−0.02	0.82	0.00	0.92	0.01	0.93
Carbon policy surprise	0.01	0.47	0.07	0.47	0.01	0.48	0.05	0.48
ΔVIX	0.00	0.80	0.02	0.82	0.00	0.99	0.00	0.99
Δ China CLI	0.02	0.25	0.08	0.27	0.02	0.12	0.09	0.15

Table A.3: Diagnostics of the Soft-Power Shock: Full Results

Panel A reports first-stage regressions of the monthly Google index of attention to the giant panda (0–100) or to China on an instrument, six lags of the dependent variable, country and month fixed effects (CT) and, in the last column, country-by-calendar-month fixed effects (CT+CM). Standard errors are clustered by month; with one instrument the cluster-robust F-statistic equals the effective F-statistic of Montiel Olea and Pflueger (2013). Instruments: births with national dose (headline), births with regional dose where usable, the count of births without dose, and hard political events (China-topic intensity times the direction of the event, +1 rapprochement, −1 conflict). The coefficient is in index points per standard deviation of dose (per birth for the count). Sample 2011–2025 (4,320 country-months, 180 months) unless stated. Panel B regresses the soft-power shock on six lags of each block of variables with country and calendar-month fixed effects; F is the Wald statistic for the block, with its asymptotic p-value and a permutation p-value from 300 draws that reshuffle the shock across years within country-by-calendar-month cells. Industrial production is available for 14 hosts, the exchange rate for 20, geopolitical risk for 20, policy uncertainty for 12; China's leading indicator (OECD CLI) is common to all. The last row is the AR(1) coefficient of the shock with country fixed effects. Panel C reports the calendar-month distribution of births. Panel D compares announced breeding attempts that produced a live birth with those that did not (pseudo-pregnancy, resorption or stillbirth, neonatal death); attention is the mean monthly index in the three months before the announcement of the attempt, raw and standardized within country; the dose of an announcement is its intensity (equation 1) with the regional source where usable, with the national source, and on the China topic; p-values are from Welch t-tests.

Panel A: First stage

Dependent	Instrument	Coef.	s.e.	F (CT)	F (CT+CM)	Event-months
Panda attention	Births, national dose	2.28	0.29	60.51	78.36	33
Panda attention	2017–2025	2.62	0.43	37.22	55.49	20
Panda attention	2004–2025	2.28	0.28	66.11	80.78	42
Panda attention	Births, regional dose	2.25	0.28	66.26	87.58	33
Panda attention	Births, count	15.63	3.46	20.37	26.41	33
Panda attention	Hard political events	−0.06	0.11	0.28	0.26	69
China attention	Births, national dose	0.03	0.09	0.10	0.59	33
China attention	Hard political events	−1.32	0.54	5.97	5.53	69

Panel B: Forecastability of the soft-power shock, 2011–2025

Block of lagged variables	Regressors	F	p (asymptotic)	p (permutation)
Attention to panda and China	12	0.81	0.64	0.69
Host industrial production growth	6	1.66	0.13	0.12
Exchange rate against CNY	6	0.47	0.83	0.92
China leading indicator, change	6	1.46	0.19	0.70
Country geopolitical risk	6	1.64	0.14	0.13
Country policy uncertainty	6	0.46	0.84	0.95
All macro and political jointly	24	2.54	0.00	0.41
Attention, FX, CLI and GPR jointly	30	2.91	0.00	0.36
AR(1) coefficient		−0.00	0.63	

Panel C: Seasonality of births

Births	Number	Share July–September	Share June–September
All births in the database	67	0.73	0.85
Baseline births, 2011–2025	33	0.73	0.85

Panel D: Announced breeding attempts, live birth versus none

Variable	N live	Mean live	N failed	Mean failed	p
Panda attention, 3 months before	22	36.26	30	32.48	0.47
standardized	22	−0.15	30	−0.10	0.84
China attention, 3 months before	22	37.39	30	33.69	0.16
standardized	22	0.23	30	−0.18	0.05
Dose of pregnancy announcement	16	0.45	10	0.83	0.63
national	16	0.20	13	0.68	0.43
China topic	16	−0.33	13	−0.03	0.31
Dose of attempt announcement	14	0.26	20	−0.03	0.16
national	17	0.19	20	−0.07	0.74
China topic	17	−0.20	20	−0.03	0.45

Table A.4: First Stage of the Soft-Power Shock: Leaving Out Hosts and Trimming the Dose

Each row re-estimates the first stage of Table A.3, Panel A, row 1: the monthly Google index of attention to the giant panda on the headline instrument (births with national dose), six lags of the index, country and month fixed effects, 2011–2025, standard errors clustered by month. "Without" drops the host from the panel. "Largest k event-months set to zero" keeps the panel and sets the k largest values of the instrument to zero. The coefficient is in index points per standard deviation of dose (per unit of $\log(1+z)$ in the last row); F is the cluster-robust first-stage statistic; the last column counts the host-months in which the instrument is nonzero.

Sample or instrument	Coef.	F	Event-months
Baseline	2.28	60.5	33
Panel A: One host left out			
Without Belgium	2.73	40.1	31
Without Canada	2.25	63.4	32
Without Germany	2.27	62.5	32
Without Spain	2.28	60.2	30
Without France	2.08	93.5	31
Without Hong Kong	2.24	60.3	32
Without Japan	2.25	59.7	28
Without South Korea	2.26	61.2	31
Without Malaysia	2.29	58.8	30
Without Netherlands	2.43	50.3	31
Without Russia	2.30	57.7	32
Without Singapore	2.22	60.9	32
Without Taiwan	2.29	59.6	31
Without United States	2.26	57.0	26
Panel B: Trimming the dose			
Largest 1 event-month set to zero	2.67	45.1	32
Largest 3 event-months set to zero	2.38	18.2	30
Largest 5 event-months set to zero	2.91	36.6	28
Largest 10 event-months set to zero	4.41	35.5	23
Dose capped at its 90th percentile	3.12	49.6	33
Dose in logs, $\log(1+z)$	12.47	51.3	32

Table A.5: Responses to One Typical Panda Birth: Alternative Instruments, Fixed Effects and Samples

Each row re-estimates the headline response of Table 3 with one change; the cells give the coefficient per typical birth, in percent (100 times the log change from the month before the announcement to h months after it), with the standard error clustered by month in parentheses. Exports and imports: 24 hosts, country and month fixed effects, control block, 2011–2025. Chinese visitors: 15 hosts including Taiwan, country, month and country×calendar-month fixed effects, months 2020-02 to 2023-01 dropped, control block with 12-month pre-growth, source-break regressors, all births kept. Headline: births with national dose, scaled to the median usable birth since 2011. Regional dose: the intensity measured in the Google series of the zoo's region where that window is usable and nationally otherwise, scaled to its own median usable birth; this row was estimated for this table with the code of the headline. Count of births: one per birth, without dose. Dose capped at its 90th percentile across the 24 hosts. Birth-month dummy: one per host-month with a birth, including births without a measurable dose. Country×year fixed effects replace the country fixed effects (month fixed effects kept). The 2004–2025 sample is estimated with the count of births only. Visitors 2011–2019 and 2017–2025: the closure months are dropped where the sample contains them. LP-IV: attention to the panda instrumented by the headline instrument, the coefficient converted to the response to a typical birth, and in brackets the Anderson–Rubin 90% interval in the same units. Blank visitor cells: the country×calendar-month row is the visitor headline itself; country×year fixed effects without country×calendar-month fixed effects leave the seasonal pattern of arrivals in the residual and are not run for visitors; LP-IV was not run for the 15-host visitor panel.

	Exports to China		Imports from China		Chinese visitors	
Specification	$h = 6$	$h = 12$	$h = 6$	$h = 12$	$h = 6$	$h = 12$
Headline	−0.28	0.16	−0.40	0.33	−0.16	1.30
	(0.82)	(0.53)	(0.71)	(0.62)	(3.11)	(2.82)
Regional dose	−0.25	0.18	−0.41	0.44	0.53	1.29
	(0.86)	(0.57)	(0.74)	(0.69)	(2.44)	(2.32)
Count of births, no dose	−2.54	0.64	−0.91	4.34	−2.33	−1.10
	(2.91)	(1.97)	(2.08)	(1.76)	(5.34)	(4.67)
Dose capped at 90th percentile	−0.10	0.56	−0.07	0.95	−0.16	1.30
	(1.42)	(0.82)	(1.07)	(0.93)	(3.11)	(2.82)
Birth-month dummy	−3.83	1.15	0.12	5.04	−2.63	0.45
	(2.65)	(1.71)	(1.93)	(1.68)	(5.84)	(4.96)
Country × calendar-month FE	−0.24	0.35	−0.33	0.26		
	(0.93)	(0.77)	(0.50)	(0.70)		
Country × year FE	−1.31	0.05	−0.58	−0.22		
	(0.72)	(0.50)	(0.59)	(0.34)		
Sample 2011–2019	−1.11	0.48	−0.67	0.06	−0.31	0.49
	(0.88)	(0.61)	(0.53)	(0.24)	(3.03)	(3.45)
Sample 2017–2025	0.23	0.53	0.63	0.86	−3.02	8.58
	(1.81)	(0.74)	(1.09)	(1.01)	(5.65)	(5.76)
Sample 2004–2025, count	−0.08	2.56	−1.76	4.12	−1.70	2.60
	(2.47)	(1.68)	(1.75)	(1.51)	(4.30)	(3.51)
LP-IV, Anderson–Rubin 90% set	−0.27	0.21	−0.31	0.41		
	[−1.65, 1.10]	[−0.74, 0.95]	[−1.62, 0.85]	[−0.74, 1.29]		

Table A.6: From an Annual Birth-Year Estimate to the Monthly Local Projection: Exports to China

The outcome is the host's exports to China in the 24 host jurisdictions (IMF and Taiwan Ministry of Finance), in percent (100 times log points). Step (1) regresses log annual exports on an indicator for a year in which at least one birth was announced, using every birth including those without a measurable dose, with host and year fixed effects, in the spirit of the annual gravity estimate of Longhi and Morelli (2025); the count row replaces the indicator with the number of births, and the November–October row uses their timing of the birth year. Step (2) adds the indicator for a birth in the following year (lead) and in the previous year (lag) to the same equation; rows using the lead drop 2025. Step (3) adds the host's export growth from year $t-2$ to $t-1$. Step (4) moves to monthly data: log exports on the number of births announced in month $t-k$ for $k=-12,\ldots,12$, with host and month fixed effects; coefficients are levels relative to months more than a year away from any birth, and the rows report the announcement month, averages over the year after and the year before, and their difference; clustering by host instead of by month raises the standard errors of the averages to about 4 and of the difference to about 2. Step (5) is the headline local projection of Section 3.4 (cumulative change from the month before, country and month fixed effects, six lags, the control block and the previous twelve months' growth) with the monthly birth count as the shock; step (6) replaces the count with the dose, the headline specification, in percent per median birth. Standard errors are clustered by host in the annual regressions (24 clusters) and by month in the monthly ones. Observations: 360 and 528 host-years in step (1), 336 and 504 in the rows with the lead, 4,320 and 6,312 host-months in step (4), and those of the headline at each horizon in steps (5) and (6).

			2011–2025		2004–2025	
Step	Regressor	Per	Coef.	SE	Coef.	SE
(1) Annual	Birth year	Birth year	8.40	4.46	5.83	4.47
	Births in year (count)	Birth year	8.05	3.66	5.94	3.62
	Birth year, Nov.–Oct.	Birth year	8.30	4.45	5.91	4.49
(2) + lead, lag	Birth in year $t+1$ (lead)	Birth year	5.36	4.92	1.80	4.43
	Birth in year t	Birth year	9.37	5.86	5.03	5.31
	Birth in year $t-1$ (lag)	Birth year	7.28	5.36	3.37	4.64
(3) + growth	Birth year, with growth control	Birth year	8.10	3.38	6.35	3.98
	Lead, with growth and lag	Birth year	4.60	4.42	1.29	4.28
(4) Monthly, levels	Announcement month, $k=0$	Birth	10.50	4.17	8.77	3.66
	Mean of $k=0,\ldots,11$	Birth	7.09	1.00	5.71	1.04
	Mean of $k=-12,\ldots,-1$	Birth	8.14	1.15	5.89	1.04
	Year after minus year before	Birth	−1.04	1.37	−0.18	1.14
(5) LP, count	Horizon $h=0$	Birth	3.01	1.51	3.42	1.26
	Horizon $h=6$	Birth	−3.83	2.65	−0.96	2.16
	Horizon $h=12$	Birth	1.15	1.71	2.55	1.49
(6) LP, dose	Horizon $h=0$	Med. birth	1.08	0.86	0.79	0.61
	Horizon $h=6$	Med. birth	−0.28	0.82	0.10	0.61
	Horizon $h=12$	Med. birth	0.16	0.53	0.29	0.46

Table A.7: Randomization Inference for the Headline Responses

Each row is a headline reduced-form local projection at horizon h (months): the coefficient on the birth shock in percent per median usable birth, with its standard error clustered by month and the asymptotic p-value. Consumer imports: imports from China in the consumer HS2 chapters. Trade outcomes: 24 hosts, country and month fixed effects and the control block; Chinese visitors: 15 hosts, country, month and country × calendar-month fixed effects, months of China's outbound closure dropped (15 of the births fall on estimation rows). Permutation tests: in each of 1,000 draws every one of the 33 usable births since 2011 keeps its host and its dose and is moved to a random month in which its host held a breeding pair (from twelve months after the pair's arrival to its return); in scheme A the calendar month is drawn from the distribution of the actual birth months, in scheme B any month of the tenure is eligible. The regression is re-estimated on the same rows with the same fixed effects and controls; p is the share of draws with an absolute coefficient at least as large as the actual one, and the null band is the 2.5th to 97.5th percentile of the placebo coefficients under scheme A, to be compared with ±1.96 times the month-clustered standard error. Wild bootstrap: restricted wild cluster bootstrap-t clustered by host (Rademacher weights, 999 draws; 24 host clusters for trade, 15 for visitors).

				p-value			Null band		Wild
Outcome	h	Coef.	SE	Asympt.	Perm. A	Perm. B	(scheme A)	1.96×SE	boot. p
Exports to China	6	−0.28	0.82	0.73	0.80	0.79	[−2.08, 2.74]	1.60	0.81
	12	0.16	0.53	0.76	0.89	0.91	[−2.30, 2.25]	1.04	0.73
Imports from China	6	−0.40	0.71	0.57	0.65	0.58	[−1.63, 1.96]	1.38	0.51
	12	0.33	0.62	0.60	0.65	0.71	[−1.32, 1.83]	1.22	0.65
Consumer imports	6	−0.79	0.90	0.38	0.41	0.37	[−2.45, 1.71]	1.76	0.43
	12	0.41	0.63	0.51	0.73	0.68	[−1.74, 2.61]	1.23	0.60
Chinese visitors	6	−0.16	3.11	0.96	0.94	0.94	[−6.85, 6.75]	6.10	0.95
	12	1.30	2.82	0.64	0.68	0.74	[−8.57, 8.30]	5.52	0.70

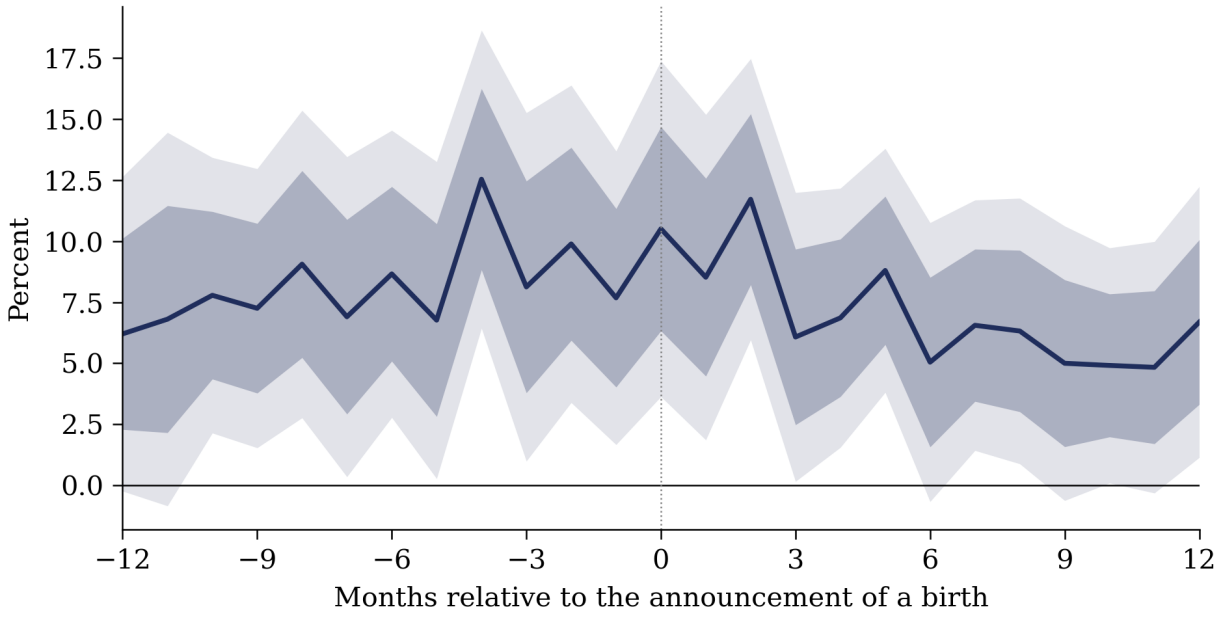


Figure A.2: **Exports to China around a Birth, Monthly Event Study in Levels.** Step (4) of Table A.6: the logarithm of the host's exports to China on the number of births announced in month $t - k$, for $k = -12, \ldots, 12$, with country and month fixed effects, every birth including those without a measurable dose, 24 hosts, 2011–2025. Each coefficient is the level of exports k months from the announcement relative to months more than a year away from any birth, in percent per birth; bands are 68 and 90 percent intervals, standard errors clustered by month.

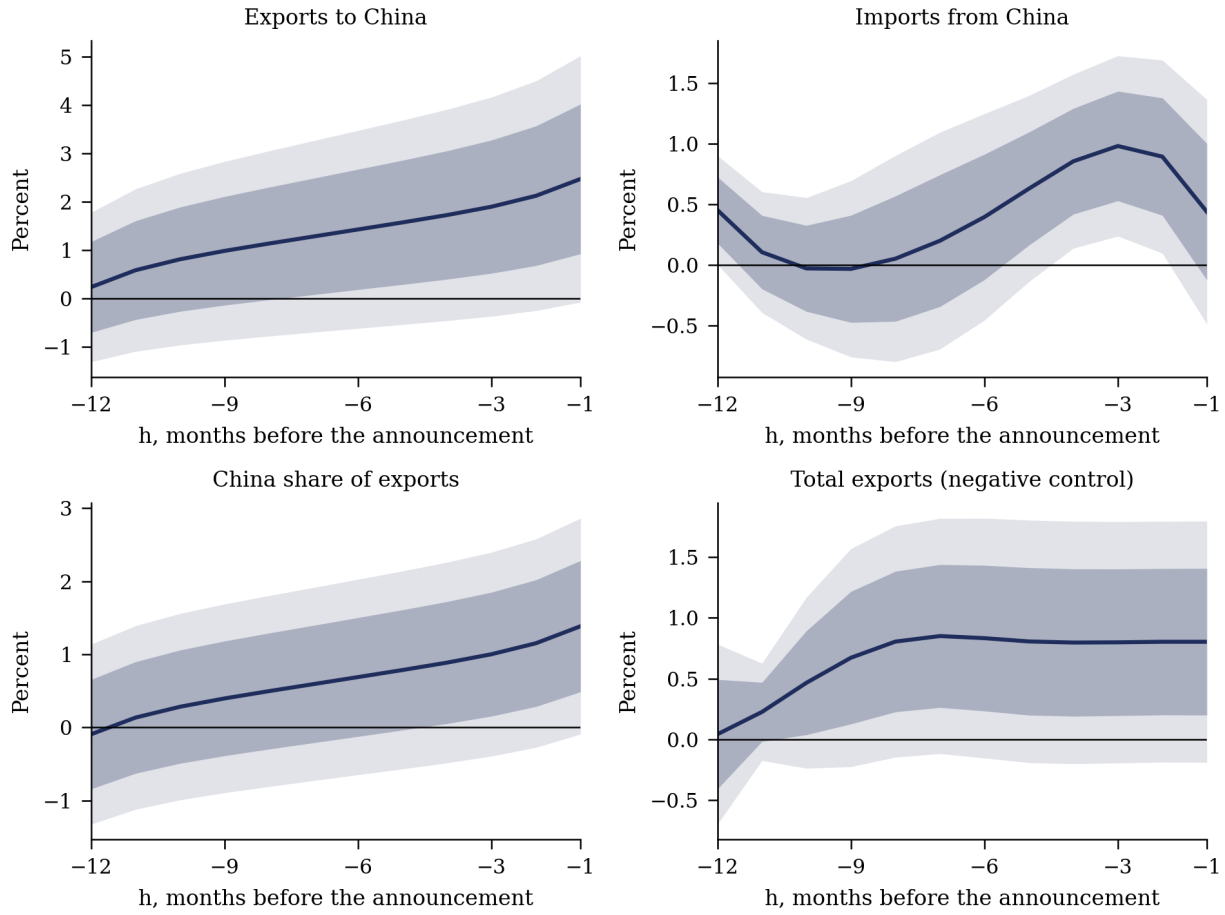


Figure A.3: **Trade before the Announcement of a Birth.** The headline regression of Figure 5 with the outcome measured before the announcement, $y_{i,t+h} - y_{i,t-13}$ for $h = -12, \ldots, -1$: the change from 13 months before the announcement to h months before it. The lags of the outcome end at $t - 13$ and the pre-growth control is dated $y_{t-14} - y_{t-26}$. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path, per typical birth, 2011–2025, country and month fixed effects; bands are 68 and 90 percent intervals, standard errors clustered by month.

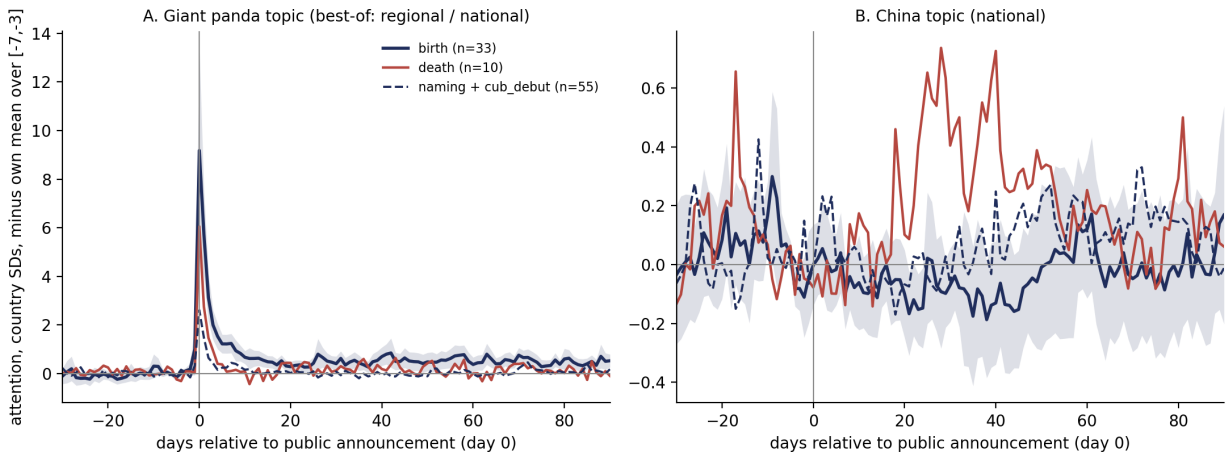


Figure A.4: **Attention around Panda Events: Daily Means, with Deaths and Namings.** Each line is the cross-event mean of daily Google attention in the host country, in units of the country's standard deviation of daily attention on non-event days, minus the event's own mean over days −7 to −3, so that the mean over days 0 to +2 equals the average intensity of equation (1). Day 0 is the first public announcement. Panel A: giant panda topic, measured in the region of the zoo where Google's regional data are usable and nationally otherwise. Panel B: China topic, national. Births (solid, $n = 33$; shaded band: mean ± 1.645 cross-event standard errors), deaths (red, $n = 10$), namings and cub debuts (dashed, $n = 55$). Baseline events announced from January 2011 to December 2025.

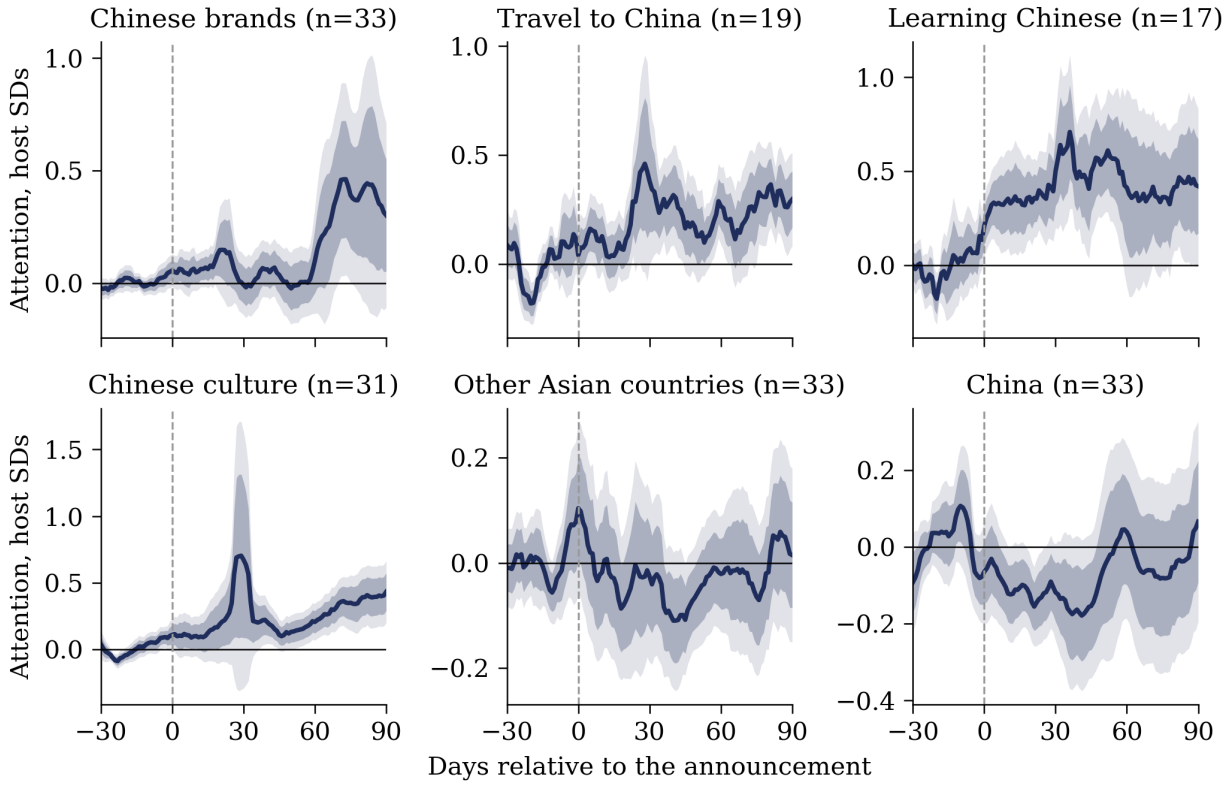


Figure A.5: **Attention to China around Panda Births, Daily Series.** Each panel plots the cross-event mean of daily Google attention in the host to one group of topics around the 33 usable births announced from 2011 to 2025, smoothed with a centered seven-day moving average. Each topic is rescaled to the host's monthly index of the same topic, divided by the host's standard deviation of that topic on non-event days and measured relative to the event's own mean over days −37 to −7; the topics of a group are averaged within each event. Brands: Huawei, Xiaomi, TikTok, SHEIN. Travel: Chengdu, the Great Wall. Learning Chinese: Mandarin. Culture: Chinese New Year, Chinese cuisine. Placebo countries: Japan, South Korea, Thailand, excluding the host's own country. China: the country topic. Day 0 is the public announcement. Dark and light bands are 68 and 90 percent intervals from cross-event standard errors of the smoothed series; n is the number of events with at least one usable topic in the group. Figure A.6 nets out the same dates in years without a birth, and Table 4 tests the differences.

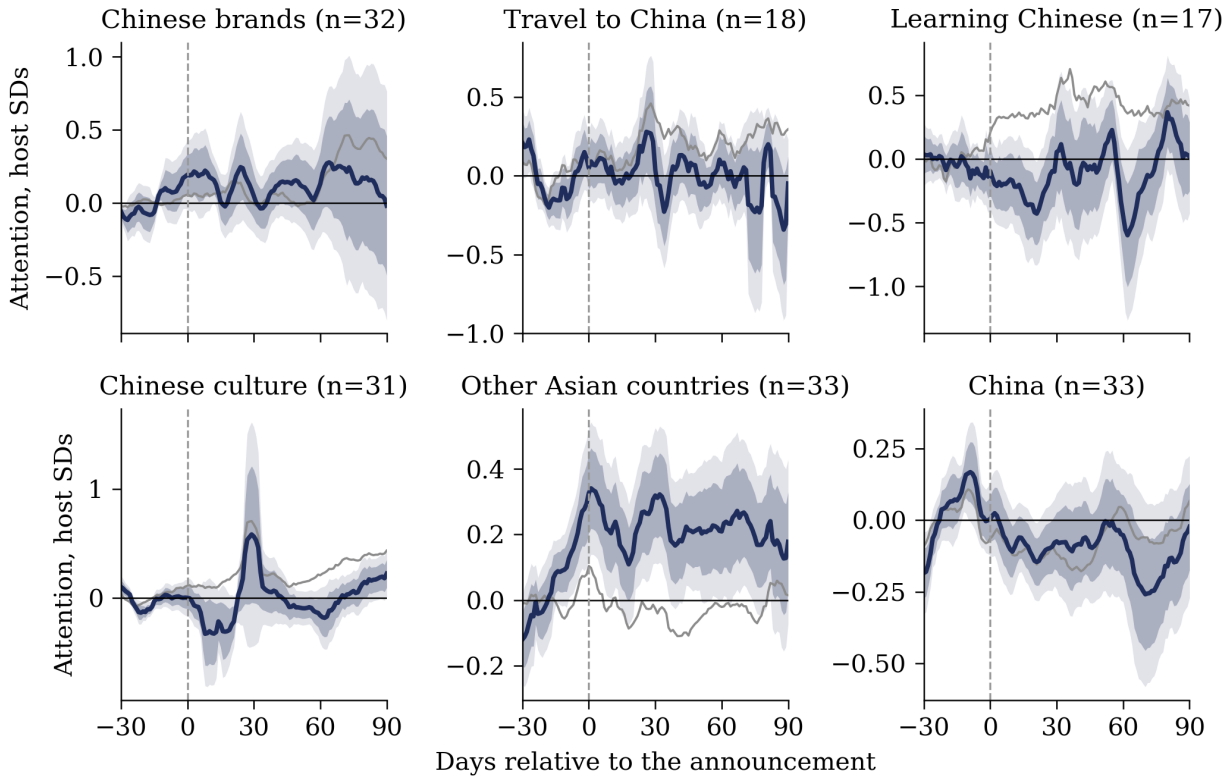


Figure A.6: **Attention to China around Panda Births, Net of the Calendar.** Each panel plots the series of Figure A.5 minus the same series for the same host on the same dates in years without a panda event: for each birth, the nearest earlier and later years from 2011 to 2025 in which the host had no panda event within 60 days of the matched day 0, which gives 60 windows for the 33 births (27 births have two, six have one). The difference is formed by topic, averaged within each group, smoothed with a centered seven-day moving average and averaged across births; the thin gray line is the unadjusted mean of Figure A.5. Scaling and standardization follow Figure A.5, with each window measured relative to its own mean over days −37 to −7. Dark and light bands are 68 and 90 percent intervals from cross-birth standard errors; n is the number of births with a matched window for the group.

Table A.8: Attention to China around Panda Births: Adjusted p-Values and Unadjusted Birth Windows

Panel A: p-values for the differences of Table 4. Holm and BH: Holm and Benjamini–Hochberg adjusted p-values over the twelve secondary tests (four groups × three windows). RI: randomization inference; in each of 2,000 draws, within every birth one of its windows (the birth window and its placebo windows) is relabelled at random as the birth window and the rest as placebo windows, and the cross-birth mean is recomputed; p is the share of draws with an absolute mean at least as large as the actual one. Panel B: birth-window means without the placebo subtraction, over all births of Appendix Figure A.5 (China 33, Chinese brands 33, Travel to China 19, Learning Chinese 17, Chinese culture 31, Other Asian countries 33); per birth and topic, daily attention is rescaled to the host's monthly index, divided by the host's standard deviation and measured relative to the window's own mean over days −37 to −7, averaged within the group and over the days of the window, on the unsmoothed daily series. SE and p: cross-birth standard error and two-sided p-value from a t distribution with $n - 1$ degrees of freedom.

	Days 0–2			Days 7–30			Days 31–90		
Panel A. Adjusted and randomization-inference p-values									
Group	Holm	BH	RI	Holm	BH	RI	Holm	BH	RI
China			0.65			0.21			0.32
Chinese brands	1.00	0.89	0.15	1.00	0.89	0.43	1.00	0.89	0.65
Travel to China	1.00	0.89	0.55	1.00	0.89	0.45	1.00	0.89	0.80
Learning Chinese	1.00	0.89	0.25	1.00	0.89	0.35	1.00	0.89	0.80
Chinese culture	1.00	0.89	0.79	1.00	0.89	0.82	1.00	0.89	0.89
Other Asian countries			0.05			0.11			0.15
Panel B. Birth window, no placebo subtraction									
Group	Est.	SE	p	Est.	SE	p	Est.	SE	p
China	−0.02	0.09	0.85	−0.12	0.06	0.07	−0.07	0.12	0.60
Chinese brands	0.07	0.05	0.20	0.07	0.06	0.25	0.19	0.15	0.21
Travel to China	0.14	0.12	0.25	0.20	0.11	0.08	0.24	0.10	0.03
Learning Chinese	0.22	0.10	0.05	0.36	0.10	<0.01	0.47	0.21	0.04
Chinese culture	0.13	0.12	0.27	0.30	0.26	0.25	0.26	0.08	<0.01
Other Asian countries	0.07	0.10	0.49	−0.04	0.08	0.64	−0.03	0.06	0.59

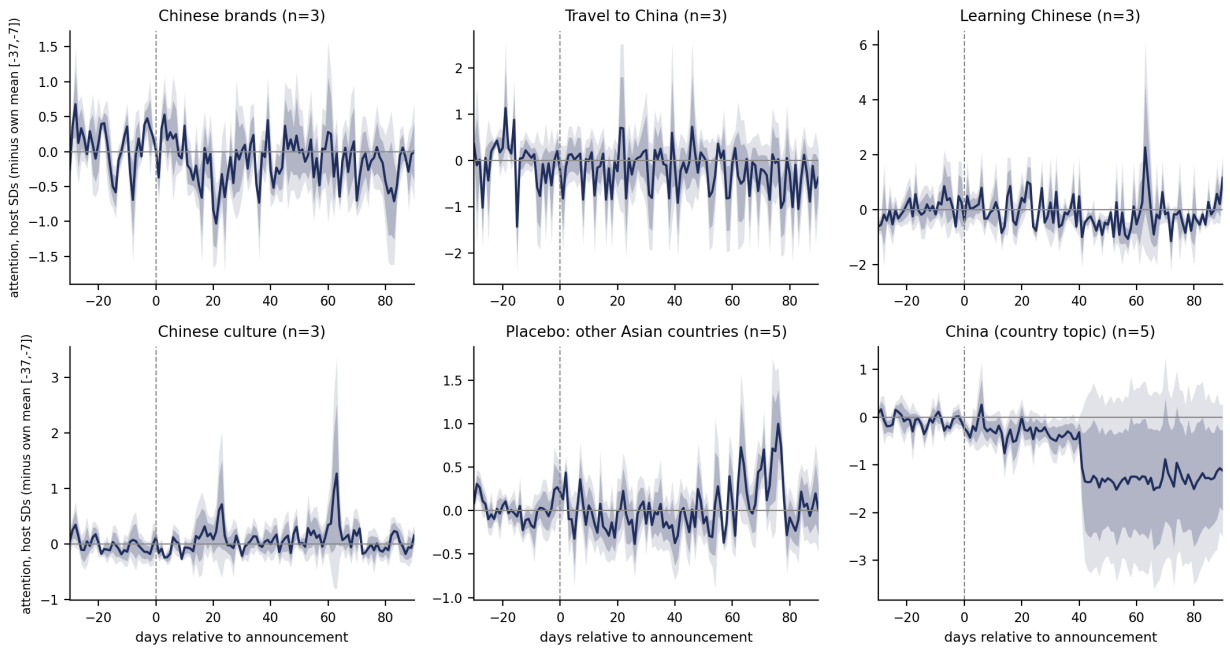


Figure A.7: **Attention to China around Placebo Animal Events.** As Figure A.5, with unsmoothed daily means, for the five animal events without a link to China that have a usable dose: Harambe (United States, 2016), Marius (Denmark, 2014), the IKEA monkey (Canada, 2012), Koshik (South Korea, 2012) and Inuka (Singapore, 2018).

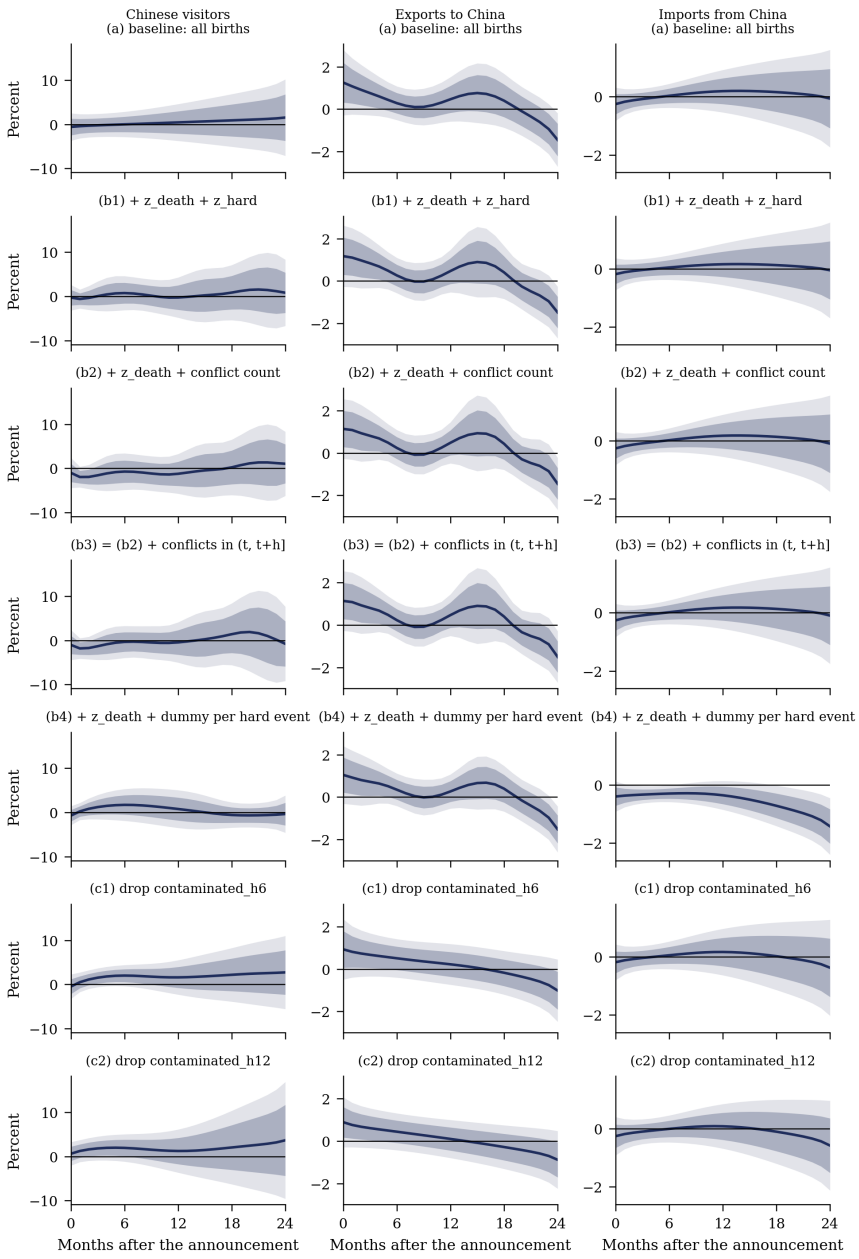


Figure A.8: **Births Followed by a Hard Political Event.** Response to one typical birth of Chinese visitors (the 15 hosts of Figure 4 without Taiwan, specification of that figure), exports to China and imports from China (24 hosts, specification of Figure 5), in percent; one column per outcome. Rows: (a) all births; (b1) the death and hard-event intensity shocks in the same equation; (b2) the death shock and the count of conflicts; (b3) (b2) plus the count of conflicts over the response window; (b4) the death shock and one window regressor for each hard event; (c1) and (c2) births followed by a conflict or a large hard event at the same host within six or twelve months set to zero. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path (paths estimated in one equation are smoothed jointly); bands are 68 and 90 percent intervals, standard errors clustered by month.

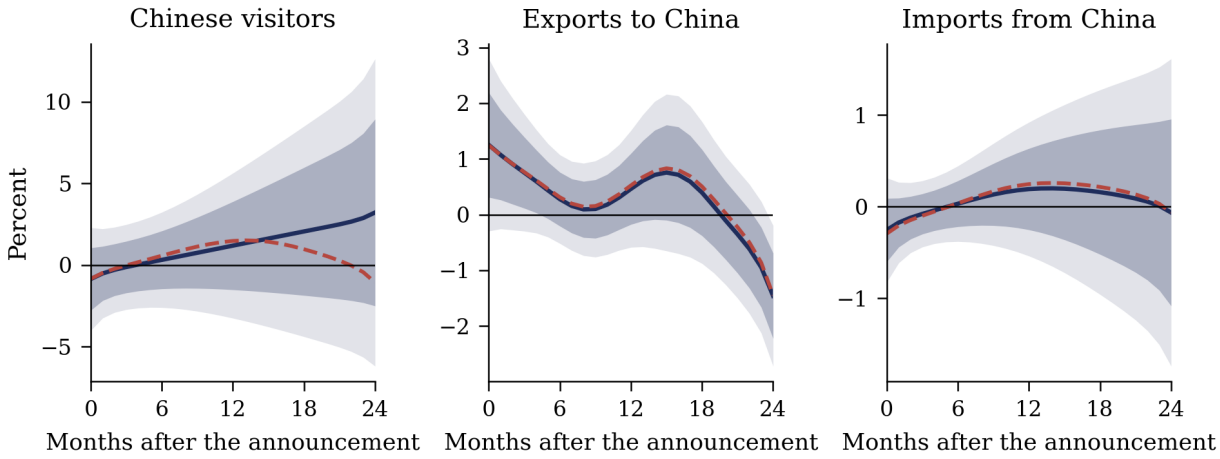


Figure A.9: **Panda Births with and without Host Policy-Regime Controls.** Navy line and bands: the headline responses to one typical birth of Chinese visitors (specification of Figure 4: 15 hosts including Taiwan, country, month and country-by-calendar-month fixed effects, months from February 2020 to January 2023 dropped, source-break regressors) and of exports to China and imports from China (specification of Figure 5: 24 hosts, country and month fixed effects), each with the control block and the outcome's growth over the previous 12 months, in percent. Red dashed line: the same projection with five host-month policy-regime variables added, each entered as its change over the response window, $D_{i,t+h} - D_{i,t-1}$: host visa easing toward Chinese visitors, China's visa-free entry for the host's citizens, free-trade agreements with China in force, Chinese group-travel bans on the host, and the suspension (−1) or resumption (+1) of direct flights; 139 dated regimes. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path; bands are 68 and 90 percent intervals, standard errors clustered by month; shock months 2011–2025.

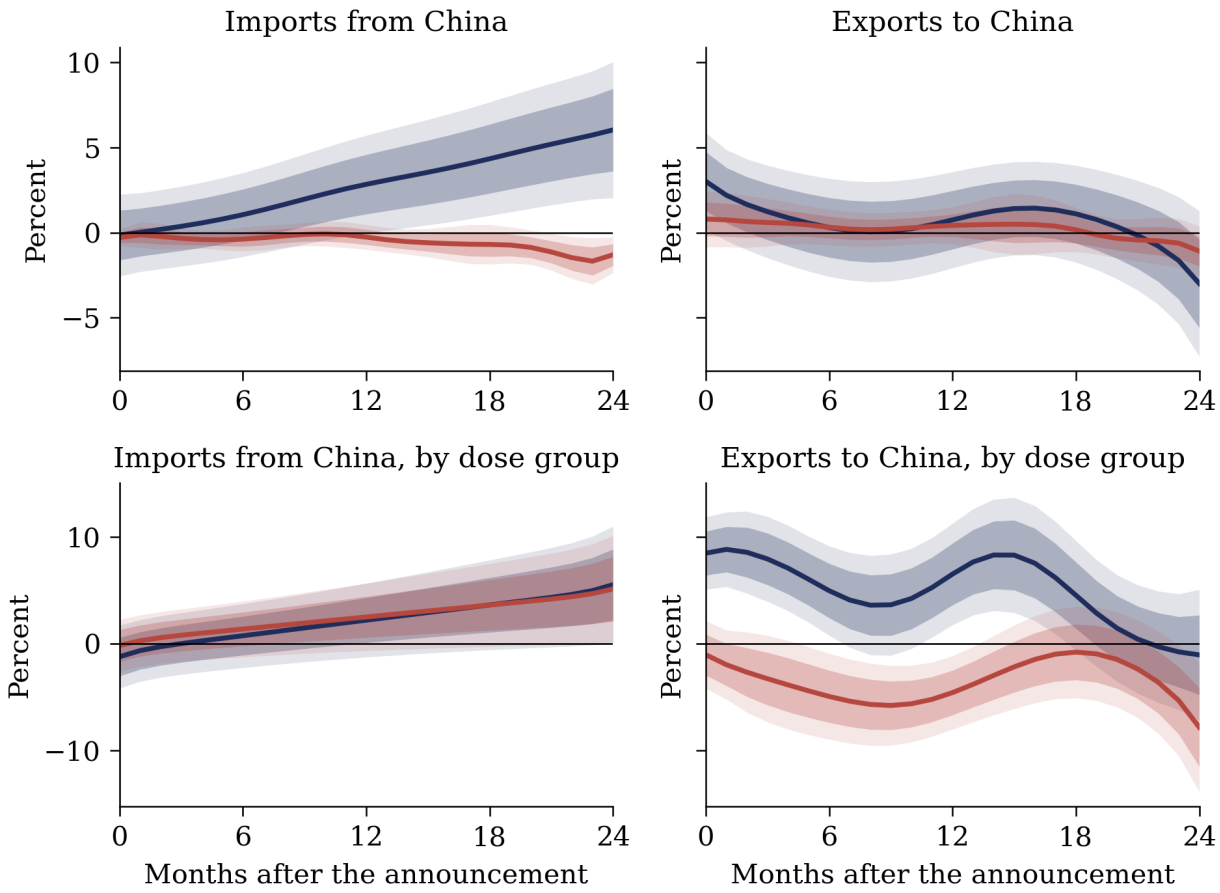


Figure A.10: **The Count of Births and the Dose.** Responses of imports from China (left) and exports to China (right), in percent (100 times the log change from the month before the announcement), in the specification of Figure 5: 24 hosts, country and month fixed effects, the control block and the outcome's growth over the previous 12 months, six lags of the outcome's change and of every shock, shock months 2011–2025. Top row: one equation with the count of usable births and the excess dose, the dose minus 3.08 times the count; navy, the response to one birth at the median usable dose (3.08 national standard deviations); red, the response to 3.08 units of attention above it. Bottom row: one equation with two counts, births with a dose at or above the median (17 since 2011, navy) and below it (16, red); percent per birth. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path (the two paths of a panel are smoothed jointly); bands are 68 and 90 percent intervals, standard errors clustered by month. With the count as the instrument, the import response at one year is 1.0 percent with a standard error of 1.2 under country-by-year fixed effects and disappears in the 2011–2019 sample; it appears in no group of HS2 chapters; it halves when Russia's trade after 2022, which the IMF estimates from China's mirror data, is excluded; and reassigning births at random within the host's tenure and the birth season gives a permutation p-value of 0.08.

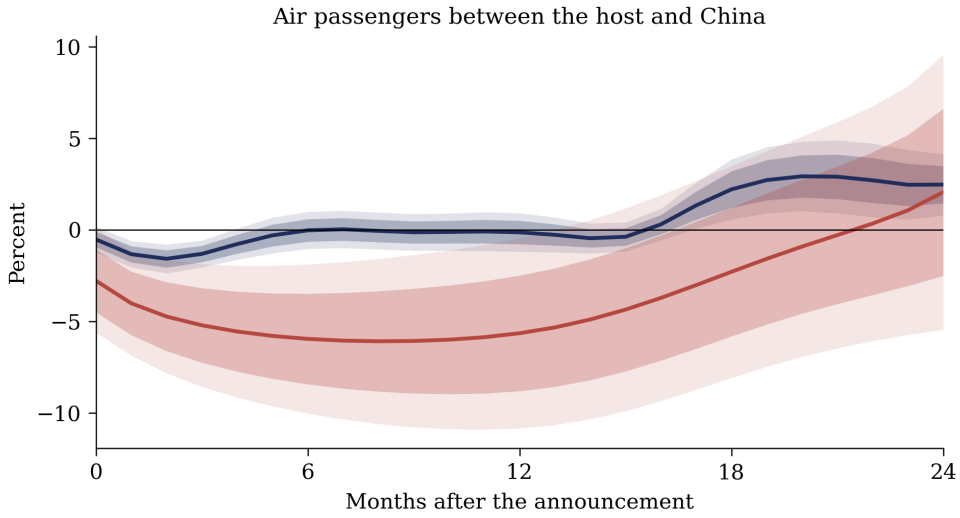


Figure A.11: **Air Passengers between the Host and Mainland China after a Birth and a Conflict.** Navy: one typical birth. Red: one conflict with China, entered as the monthly count of conflict-direction hard events in the host. Both shocks, a death shock and six lags of each enter the same local projection; the death path is not drawn. The outcome is passengers on direct flights between the host and mainland China in both directions (percent, 100 times the log change from the month before the event). Twelve hosts with at least 120 months of data: the eight EU hosts (Eurostat), the United States (Department of Transportation, T-100), Australia (BITRE city pairs), the United Kingdom (Civil Aviation Authority) and Taiwan (Civil Aviation Administration, cross-strait routes from June 2011); 14 births identify the path. Country, month and country-by-calendar-month fixed effects; six lags of the outcome's monthly change and of the shocks; the control block and the outcome's growth over the previous 12 months; months from February 2020 to January 2023 dropped. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path (the paths are smoothed jointly); bands are 68 and 90 percent intervals, standard errors clustered by month; shock months 2011–2025.

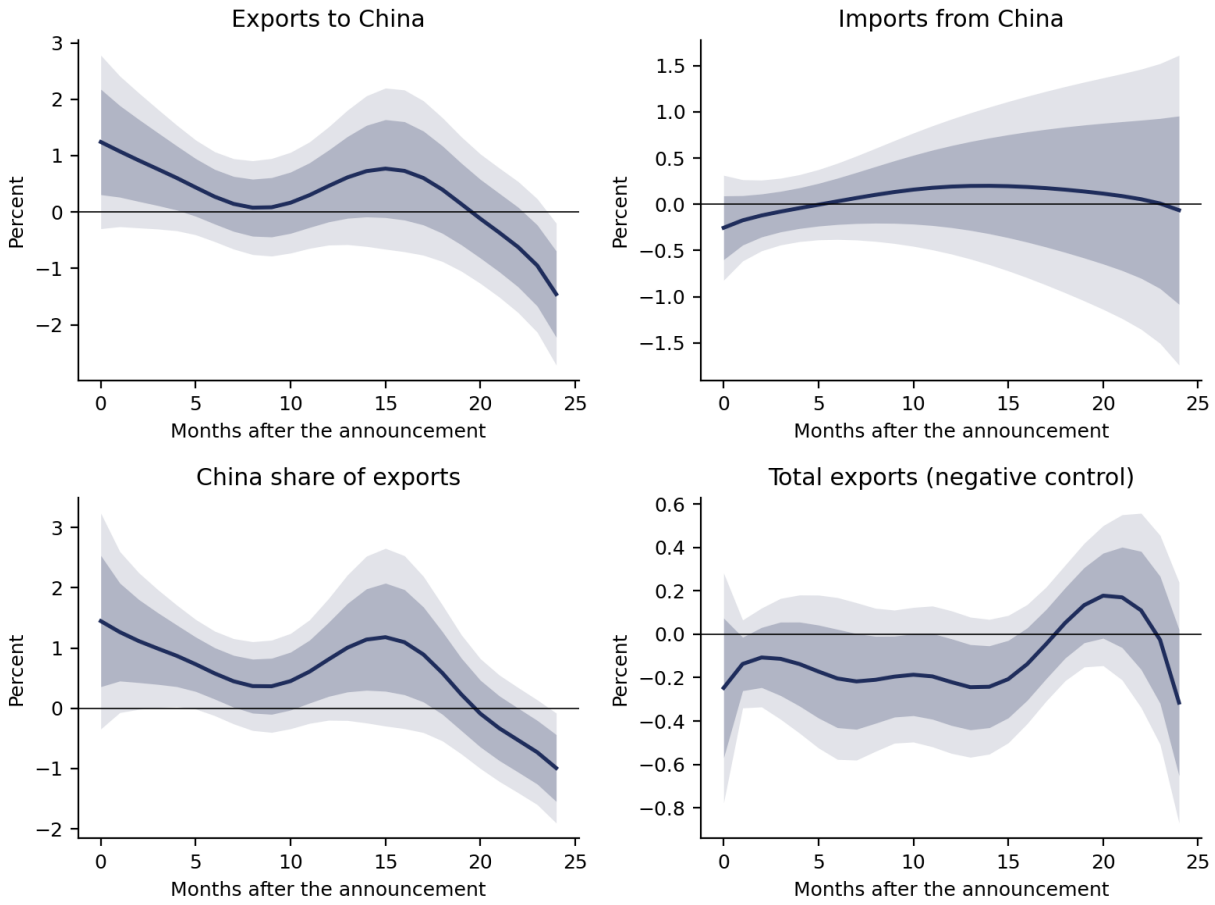


Figure A.12: **Bilateral Trade with China after a Panda Birth: All Four Series.** Each panel plots the response to one typical birth of the host's exports to China, its imports from China, China's share of its exports and its total exports (percent, 100 times the log change from the month before the announcement). Country and month fixed effects; six lags of the outcome's monthly change and of the shock; lags of the exchange rate against the renminbi, country geopolitical risk and policy uncertainty, China's leading indicator, host total exports and the other bilateral flow with China; the outcome's growth over the previous 12 months. Host-reported trade in US dollars (IMF; Taiwan's Ministry of Finance), 24 hosts, shock months 2011–2025. Smooth local projection (Barnichon and Brownlees 2019), penalty chosen by cross-validation. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month.

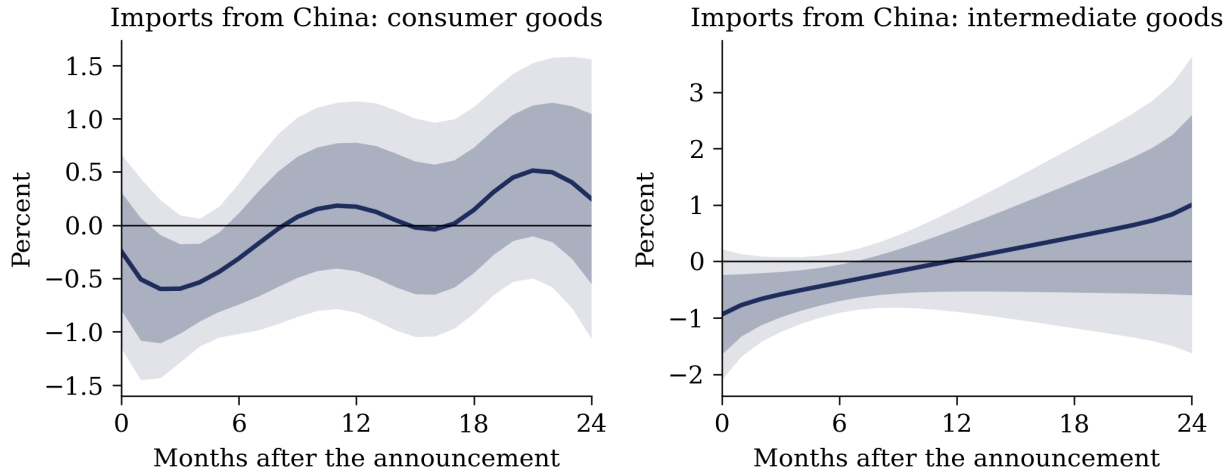


Figure A.13: **Imports from China by Type of Good after a Panda Birth.** Each panel plots the response to one typical birth of the host's imports from China in one group of two-digit HS chapters (percent). Consumer goods: chapters 02–04, 07–09, 16–22, 33, 61–64 and 95; intermediate goods: 26, 28–32, 34–39 and 72–81. Specification as in Figure 5. UN Comtrade, host-reported; Eurostat Comext for the eight EU hosts; China's mirror records for Taiwan; US dollars. Smooth local projection (Barnichon and Brownlees 2019) with one rule for both panels: the penalty is the larger of the cross-validated penalty and the penalty at which the effective degrees of freedom equal six. Dark and light bands are 68 and 90 percent confidence intervals, with standard errors clustered by month; shock months 2011–2025.

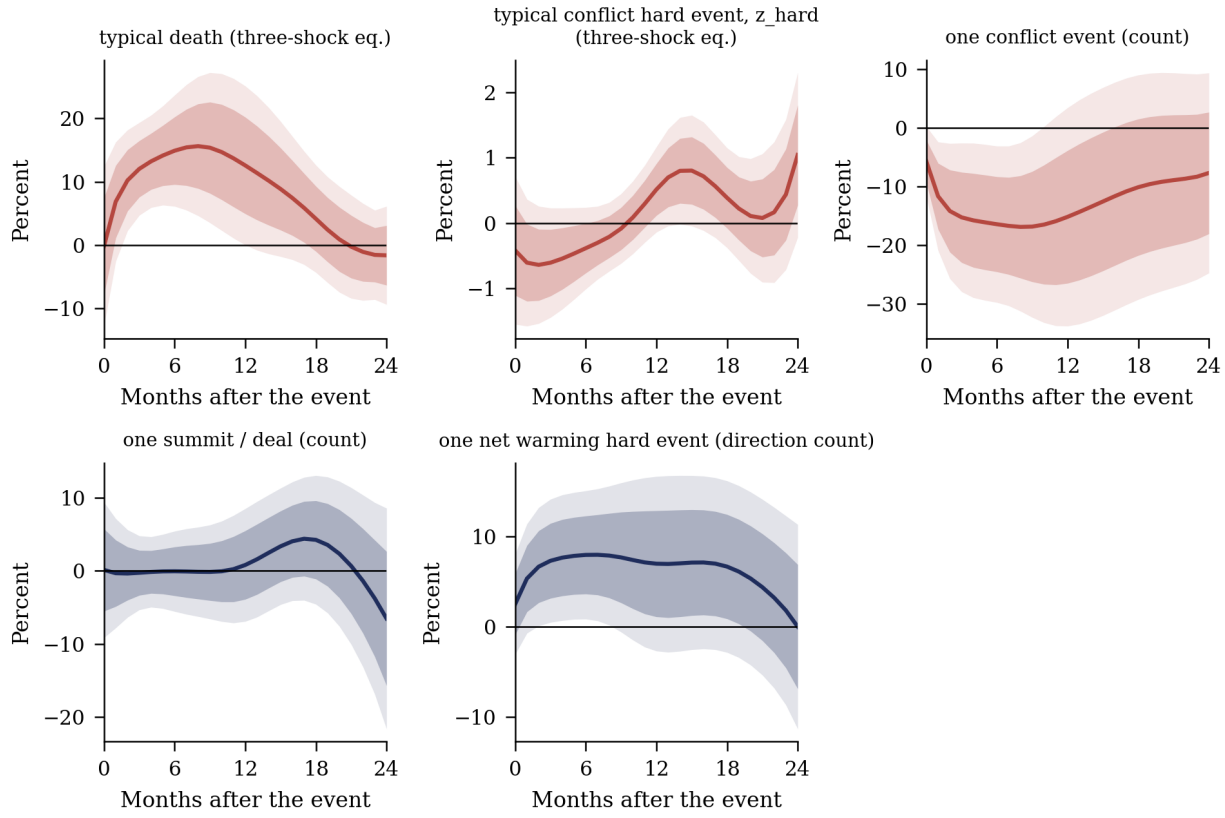


Figure A.14: **Chinese Visitors after Other Shocks.** The 15 tourism hosts of Figure 4 without Taiwan, sample and fixed effects of that figure, no control block. Panels 1–2: a typical death and a conflict of median China-topic intensity, estimated together with the soft-power shock. Panels 3–4: the counts of conflicts and of summits, entered together, per event. Panel 5: the net count of summits minus conflicts. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path (paths estimated in one equation are smoothed jointly); bands are 68 and 90 percent intervals, standard errors clustered by month.

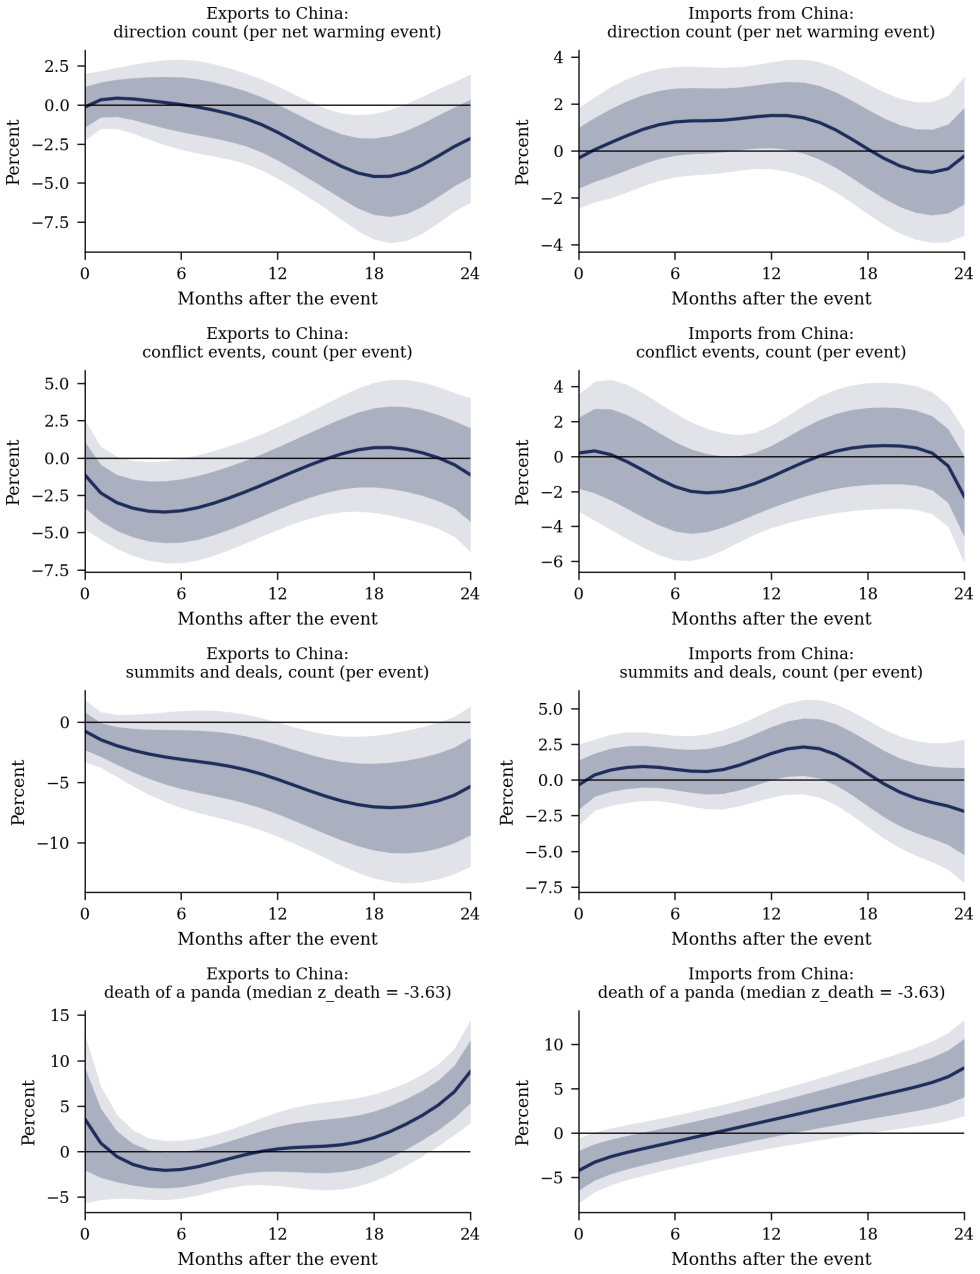


Figure A.15: **Trade after Hard Events Counted without Intensity, and after Deaths.** Left column: exports to China; right column: imports from China, percent. The first three rows replace the intensity of hard events by counts of the month's events (country and month fixed effects, no control block, 2011–2025): the net count of summits minus conflicts, and the counts of conflicts and of summits entered together. The last row is the death path from the equation of Figure 6, Panel A. Smooth local projections (Barnichon and Brownlees 2019), with the penalty chosen by cross-validation and capped at six effective degrees of freedom per path (paths estimated in one equation are smoothed jointly); bands are 68 and 90 percent intervals, standard errors clustered by month.

Appendix B Panda Diplomacy: Institutions and History

This appendix describes the institutions behind the pandas in our sample. It covers how the terms of panda diplomacy changed over time, the legal form of the current agreements, the Chinese institutions involved, the hosts, the timing of loans relative to state visits, and the latest cycle of returns and renewals.

Buckingham et al. (2013) distinguish three phases, summarized in Table B.1. In the first, during the Mao era, China gave pandas to friendly states as gifts. The best-known gift followed President Nixon's visit to Beijing in February 1972. Premier Zhou Enlai gave two pandas to the American people, and they arrived at the Smithsonian's National Zoo in April 1972 (Smithsonian's National Zoo and Conservation Biology Institute 2025). In the second phase, after Deng Xiaoping's rise in 1978, gifts became loans under a lease model based on financial transactions. In the third, which Buckingham et al. (2013) date from the aftermath of the 2008 Sichuan earthquake, long-term loans have been associated with countries that supply China with resources and technology. Hartig (2013) places the practice within China's public diplomacy and notes that it combines conservation and research with a political purpose.

The current agreements are contracts between a host zoo and a Chinese counterpart for cooperative research and breeding. The Smithsonian's agreement of 2000 ran for ten years in exchange for US$10 million and has been renewed three times. Its latest version, signed in 2024 with the China Wildlife Conservation Association, runs to April 2034 (Smithsonian's National Zoo and Conservation Biology Institute 2025). Under that agreement the zoo pays an annual fee of US$1 million to support research and conservation in China, which the Smithsonian describes as the practice of all zoos outside China that keep pandas, and the two sides conduct joint research projects in China and at the zoo. Cubs born abroad return to China. Under the Smithsonian's original agreement a cub was to be sent to a breeding center in China after its second birthday, and under the 2010 renewal after four years (Smithsonian's National Zoo and Conservation Biology Institute 2025). Our database

records 33 departures of cubs to China and 21 returns of adult pandas at the end of a loan (Table 1).

On the Chinese side, the counterparty of the American agreements is the China Wildlife Conservation Association (Smithsonian's National Zoo and Conservation Biology Institute 2025; Office of the Governor of California 2024), and the National Forestry and Grassland Administration, formerly the State Forestry Administration, oversees the program and reports on it (China Daily 2023). The pandas themselves come from the breeding centers in Sichuan, in particular those of the China Conservation and Research Center for the Giant Panda. The first pandas of the San Diego Zoo came from Wolong in 1996 (NBC 7 San Diego 2024), the Smithsonian's first surviving cub went to the Bifengxia center in 2010 and its current pair was born at the Center (Smithsonian's National Zoo and Conservation Biology Institute 2025), and the pair sent to Hong Kong in 2024 left from Dujiangyan (Hong Kong Government News 2024). In 2023 the forestry administration reported that China had cooperated with 26 institutions in 20 countries over the years, with 41 litters and 68 cubs born abroad, that 63 pandas were then living in 19 countries, and that its field assessment that year covered 23 institutions in 19 of them (China Daily 2023).

Appendix Table A.1 lists the 24 host jurisdictions in our sample, among them Hong Kong, Macau and Taiwan, with their events and the years of their first and last events. The United States and Japan have the longest histories and the most births. Several European hosts joined in the 2010s, namely Belgium, Denmark, the Netherlands and Finland. Mexico's pandas descend from the gift era, and its births, from 1980 to 1990, fall before our sample.

Decisions about loans are taken by governments and are often announced during state visits. In our database, 18 of the 25 loan announcements and 6 of the 17 agreements fall within seven days of a state visit, summit or major agreement (Table 1). Two examples show the pattern. The agreement for Edinburgh Zoo was signed in January 2011 during an official visit to the United Kingdom by Vice Premier Li Keqiang, and the pandas arrived in

December 2011. The agreement for Pairi Daiza in Belgium was signed in September 2013 at a meeting between the Belgian and Chinese prime ministers, and President Xi Jinping and King Philippe visited the newly arrived pandas there on March 31, 2014, during Xi's state visit to Belgium. Both cases are recorded, with their sources, in our event database (Section 3.1). Buckingham et al. (2013) document further coincidences between loans and trade agreements. This bundling is the reason the paper does not use loans, arrivals or returns as shocks.

Between 2023 and 2024 several American loans ended. According to our event database, the Memphis Zoo's last panda left in April 2023 and Zoo Atlanta's pandas in October 2024. The three pandas of the Smithsonian's National Zoo left on November 8, 2023 (Smithsonian's National Zoo and Conservation Biology Institute 2025). A week after the Smithsonian departures, President Xi said in San Francisco that China was "ready to continue our cooperation with the United States on panda conservation" (Xi 2023). New agreements followed. The San Diego Zoo Wildlife Alliance signed a cooperative agreement with the China Wildlife Conservation Association in February 2024 (Office of the Governor of California 2024), and a new pair arrived in June 2024 (NBC 7 San Diego 2024). The Smithsonian's renewed agreement brought a new pair to Washington on October 15, 2024, which went on public view on January 24, 2025 (Smithsonian's National Zoo and Conservation Biology Institute 2025). In Hong Kong, two pandas given by the central government to the special administrative region arrived on September 26, 2024 (Hong Kong Government News 2024). Zoos in Madrid, Adelaide and Vienna also received new pairs in 2024 and 2025, which our database records with their sources.

Table B.1: The Three Eras of Panda Diplomacy

Phases as defined by Buckingham et al. (2013); examples and terms from Smithsonian's National Zoo and Conservation Biology Institute (2025). The third phase is dated by Buckingham et al. (2013) from the aftermath of the 2008 Sichuan earthquake; the long-term agreements it builds on predate it.

Phase	Form	Example
1. Mao era	State gifts to friendly countries	Two pandas given to the American people after the 1972 Nixon visit
2. After 1978	Loans for a fee under a lease model	–
3. After 2008	Long-term research and breeding agreements	Smithsonian: ten years from 2000 for US$10 million; US$1 million a year to the China Wildlife Conservation Association; renewed to 2034