Semiconductor Innovation, Durables Consumption, and Cross Section of Stock Returns

Amir Hosseini[†]

May 2026

Abstract

I study the macroeconomic and asset-pricing implications of semiconductor innovation using a monthly innovation index constructed from granted semiconductor patents. The index aggregates patent-level textual novelty, measured from sentence-level semantic similarity using Sentence-BERT (SBERT) language model. I find that a rise in semiconductor innovations is followed by a lower producer price index for chips, which is subsequently passed through to prices of chip-intensive durable goods, most notably motor vehicles, within a few months. An expectation of lower prices of big-ticket durables leads to a drop in demand for motor vehicles and a slowdown in growth in personal consumption expenditure. I support this using consumer survey data on expected future prices of durables and vehicles, together with aggregate retail sales data for new vehicles. Motivated by these findings, I examine the price of risk associated with semiconductor innovation shocks and find that stocks with more negative exposure earn higher returns, as their payoffs are lower in states of lower consumption following a rise in innovation activity.

JEL classification: O33, O47, E44, E21, E32
Keywords: patent novelty, semiconductor innovation, process innovation, durables consumption, motor vehicles

[†]A. Hosseini is a postdoctoral research fellow at the Gillmore Center for Financial Technology, Warwick Business School (WBS), University of Warwick. Email: amir.hosseini@wbs.ac.uk

1. Introduction

The role of technological innovations as a key driver of productivity and long-term economic growth has been documented in both theoretical and empirical research (see Griliches and Lichtenberg, 1984; Crespi et al., 2008; Kung and Schmid, 2015; Kogan et al., 2017; Acemoglu et al., 2018). In contrast, technological change can be contractionary in the short run, as positive technology shocks often lead to temporary declines in real activity before longer-term productivity gains materialize (see Francis and Ramey, 2005; Basu, Fernald and Kimball, 2006). Among the most consequential innovations are those in semiconductors, which have played a pivotal role in productivity growth since the industry's emergence in the latter half of the twentieth century (Bloom, Sadun and Reenen, 2012).

In this paper, I examine the short-run effects of major semiconductor innovations on the cost of chip production and the prices of durable goods, and study how these changes affect consumer spending and aggregate consumption dynamics. I then investigate the asset pricing implications of the resulting consumption movements and show that exposure to persistent semiconductor innovation shocks constitutes a systematic source of consumption risk.

One channel through which innovation can affect durable goods prices is reductions in production costs (Abernathy, 1975; Klepper, 1996; Adner and Levinthal, 2001; Tham et al., 2026). Consistent with this mechanism, I document that major semiconductor innovations lower the cost of chip production.[1] These cost reductions are subsequently transmitted to downstream industries and within a few months, declines in input chip prices are passed through to the prices of durable goods, particularly motor vehicles. In contrast, I do not find similar effects for electronics prices, suggesting that the impact of

[1]The literature distinguishes between two main channels through which innovation affects durable goods prices. Product innovation refers to developments that improve product performance or functionality, whereas process innovation refers to developments that reduce production costs without altering product functionality (Griliches, 1961; Abernathy, 1975; Klepper, 1996; Adner and Levinthal, 2001).

semiconductor innovation in this sector operates primarily through improvements in chip performance and capacity rather than through reductions in production costs.

I next examine whether semiconductor innovations affect the timing of durable goods consumption, particularly for big-ticket items that require substantial household commitment. Using monthly survey data on large household durables and vehicles, I show that consumers revise their price expectations following innovation episodes, as reflected in an increase in the share of respondents who report that current prices are too high. Complementing this evidence, my analysis of monthly aggregate retail sales for new vehicles indicates a decline in demand for new vehicles after semiconductor innovations. This channel represents the primary way through which semiconductor innovation affects aggregate consumption in the short run.

I further observe a persistent growth in semiconductor innovations over the years panning multiple decades, highlighting the importance of its compounding effect on aggregate consumption dynamics and the asset returns. As contractions in durable consumption raise marginal utility, assets whose payoffs covary positively with semiconductor innovation shocks act as effective hedges against consumption growth risk and therefore earn lower expected returns. I provide direct empirical evidence in support of this using a two-stage Fama-MacBeth regression framework (see Fama and MacBeth, 1973), by showing that stock-level exposure to semiconductor innovation shocks is a predictor of lower expected returns in the cross-section of stock returns. This finding reinforces prior research highlighting the central role of durables consumption in explaining return premia, given its greater cyclicality and volatility relative to non-durables consumption (see Yogo, 2006; Gomes, Kogan and Yogo, 2009).[2]

Central to my analysis is the semiconductor innovation index constructed from the textual content

[2] Yogo (2006) argues that the standard Consumption Capital Asset Pricing Model (CCAPM) fails to account for the cross-section of returns when relying solely on nondurable consumption growth (see Mankiw and Shapiro, 1984; Breeden, Gibbons and Litzenberger, 1989). By incorporating the strongly procyclical nature of durables, the model better captures variation in marginal utility over the business cycle.

of granted patents in semiconductor-related technologies over the past five decades. The semiconductor innovation index is defined as the ratio of "truly novel" patents to total patents granted in a given period. I classify a patent as truly novel if its score in textual novelty lies above the 90th percentile of the sample distribution. To measure patent-level novelty I employ embedding techniques, which have recently gained traction for their superior performance in capturing semantic content in textual data (see Graeber, Roth and Zimmermann, 2024; Breitung and Müller, 2025). In particular, I use Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) to map the textual content of patent abstracts into a numerical vector in high dimensional space (embedding space) to calculate the cosine similarity between patent embeddings.[3] I show that this measure improves on the benchmark novelty measure in the literature (see Kelly et al., 2021) in predicting patent quality ex-ante, as proxied by future citations.

I use my monthly innovation index to examine the macroeconomic effects of major breakthroughs in semiconductor technologies. My empirical strategy focuses on two main mechanisms: an economic channel and a consumer behavior channel. The economic channel captures how semiconductor innovations reduce the cost of chip production and place downward pressure on motor vehicle prices in the short run. To test this channel, I rely primarily on monthly macroeconomic indicators, including the Producer Price Index for Semiconductor Components (PPI^{semi}) – a proxy for the cost of chip production – and the Consumer Price Index for new vehicles ($CPI^{new\,vehicle}$). The consumer behavior channel examines how households adjust their expectations and spending decisions in response to semiconductor innovation. If consumers anticipate that technological progress will lower the future price of durable goods, they may postpone purchases of vehicles and other big-ticket items. To test this channel, I use monthly survey data on consumer price expectations from the University of Michigan, together with data on aggregate retail

[3] SBERT modifies the BERT network to derive semantically meaningful sentence embeddings that can be compared using cosine-similarity. The architecture of SBERT allows for the derivation of fixed-sized sentence embeddings that are designed such that semantically similar sentences are close in embedding space (Reimers and Gurevych, 2019).

sales of vehicles and total vehicle units sold.

My findings show that PPI^{semi} declines following semiconductor innovation within two months. This suggests that, on average, the effects of major semiconductor innovations on the prices of chips sold to downstream industries materialize roughly two months after the associated patents are granted. I further demonstrate that $CPI^{new\,vehicle}$ typically responds to shocks in chip prices with a lag of about seven months, implying that semiconductor innovations should predict lower vehicle prices within roughly nine months. I test this and confirm that my innovation index is a predictor of lower vehicle prices within nine months, indicating that the effect is through reductions in input chip prices used for manufacturing motor vehicles.

Similarly, my analysis of monthly consumer survey data shows that, approximately two months after semiconductor innovations, the share of respondents reporting that it is a bad time to purchase large household durables or vehicles increases. Further analysis indicates that this shift is driven by consumers' perceptions that current prices are too high. I also find that these revisions in price expectations lead to a decline in demand for new vehicles. Using monthly data on new vehicle sales, I show that a one standard deviation increase in semiconductor innovation growth reduces car sales by up to 12.5 percent of their standard deviation within two months. This represents an economically meaningful effect in the context of aggregate consumption movements, given that motor vehicle sales account for an average of 7.7 percent of total monthly durable goods consumption.

The innovation-driven slowdown in durables consumption is particularly relevant for my analysis of the price of risk associated with exposure to semiconductor innovations, as it directly connects my findings to Yogo (2006), highlighting the role of durables consumption risk in explaining the cross-section of stock returns. I hypothesize a negative price of risk associated with exposure to innovation shocks, as assets with higher loadings on these shocks deliver higher payoff precisely in states where

aggregate consumption falls and the marginal utility of consumption is elevated. As such, they provide valuable consumption insurance, and equilibrium pricing requires that they earn lower expected returns.[4]

I estimate the asset-pricing implications of semiconductor innovation shocks using a Fama–MacBeth two-stage framework (Fama and MacBeth, 1973). Specifically, I examine whether cross-sectional differences in stock returns are explained by firms' exposure to unexpected monthly shocks in semiconductor innovation, while controlling for standard asset-pricing factors. I find that the estimated price of risk associated with semiconductor innovation exposure is negative and statistically significant on average, and that this pattern persists across a majority of sample years, indicating that semiconductor innovation risk has been a recurring feature of the cross section of stock returns over the past several decades.

This study is related to the literature that examines the relationship between innovation, proxied by patenting activity, and stock market returns (see Pakes, 1985; Nicholas, 2008; Hirshleifer, Hsu and Li, 2013, 2018; Haddad, Ho and Loualiche, 2022). Prior studies in this literature focus on capturing the rate of innovation or innovation quality, typically measured through patent counts or citation based indicators. In contrast, my approach measures innovation novelty — or the degree of dissimilarity — directly from the textual content of patents, following the approach proposed by Kelly et al. (2021).

To assess similarity, I depart from traditional methods that rely on word count and instead employ embedding techniques that represent entire sentences as dense numerical vectors, where semantically similar texts occupy nearby positions in the embedding space. This approach is the central methodological contribution of my study, as it extends the measurement of textual similarity beyond word count vectors (see Hoberg and Phillips, 2010, 2016; Cohen, Malloy and Nguyen, 2020; Kelly et al., 2021).[5]

[4] In equilibrium asset pricing, an asset's expected return is determined by its covariance with the representative agent's stochastic discount factor (SDF), which under standard consumption-based models is closely linked to the intertemporal marginal rate of substitution in consumption Lucas Jr (1978); Breeden (1979).

[5] The conventional approach relies on representing documents as vectors of word frequencies and computing cosine similarity between them.

My results on the effect of innovation on durable goods prices contribute to the literature examining the mechanisms through which innovation affects the prices of durables (see Griliches, 1961; Abernathy, 1975; Klepper, 1996; Adner and Levinthal, 2001). In particular, my findings on the effect of semiconductor innovation on motor vehicle prices provide evidence in support of the production-cost channel, whereby innovation lowers downstream prices by reducing input costs.

This study also contributes to the literature that views innovation as a source of systematic risk, either through its effects on productivity growth and long-run consumption growth risk (see Kung and Schmid, 2015) or through creative destruction that redistributes future wealth toward innovators (see Kogan, Papanikolaou and Stoffman, 2020). My study, however, provides evidence for a distinct channel, where innovation is a source of systematic risk through its effect on contemporaneous consumption.

The remainder of the paper is organized as follows. Section 2 details the methodology used to measure patent novelty and construct the semiconductor innovation index. Section 3 documents the effect of semiconductor innovation on prices of chips and chip-intensive durables. Section 4 investigates consumption behavior and timing following semiconductor innovations. Section 5 presents the portfolio sorting exercise to estimate the price of risk associated with exposure to semiconductor innovation shocks. Section 6 concludes.

2. Data and Methodology

I obtain my patent data from the United States Patent and Trademark Office (USPTO). My analysis focuses on semiconductor-device patents, identified using the Cooperative Patent Classification (CPC) 4-digit subclass code. Figure 1, Panel A reports total unique patent counts by CPC subclass. H01L ranks fourth overall among 672 subclasses; notably, two of the three subclasses that surpass it–G06F (digital computing) and H04L (digital communications)–are themselves downstream of semiconductor technology, as advances in transistor density define the performance frontier of computing and communications

hardware.[6] Panel B plots the monthly share of H01L patents in total grants. This share fluctuates between approximately 10% and 30% over the sample period, stabilizing near 20% from 2010 onward. Taken together, the evidence suggests that semiconductor innovation constitutes a substantial and persistent share of U.S. patenting activity.

[Insert Fig. 1 near here]

The dataset includes 472,746 semiconductor-device patents granted to both US and international assignees between January 1976 and December 2023. Assignees refer to the individuals or organizations holding legal ownership of the patent rights. Because I use USPTO records, my sample is limited to assignees who filed their patents with the USPTO, regardless of their country of origin. Figure 2 shows the distribution of granted semiconductor patents across the top five countries, based on the location of the assignee organization.

[Insert Fig. 2 near here]

Figure 3 presents the cumulative number of semiconductor patents granted by country of assignee from 1976 to 2023 for the top patenting nations. The U.S. has maintained a dominant position throughout the sample period, accounting for the largest share of granted patents, followed by Japan, Taiwan, and South Korea. The figure also highlights the rapid acceleration of U.S. and East Asian patenting activity beginning in the mid-1990s, pointing to the globalization of semiconductor supply chains. In addition, China's cumulative patenting activity while modest until the early 2010s, after has increased markedly, reflecting the country's concerted efforts to expand domestic semiconductor capabilities. Overall, the

[6] Y10T is not a substantive technology class. It is an administrative category created to house legacy U.S. patents that lacked a direct mapping into the Cooperative Patent Classification system when the USPTO transitioned from its domestic classification scheme in 2013.

figure underscores the persistent technological leadership of the United States alongside the rising innovation capacity of East Asian economies, illustrating the increasingly globalized nature of semiconductor innovation and patenting over the past five decades.

[Insert Fig. 3 near here]

In the next section, I use the textual content of semiconductor patents to construct a pairwise similarity measure that identifies patents least similar (novel) to prior ones, serving as a proxy for the novelty of their innovation. To develop this measure, I rely exclusively on patent abstracts, as they offer a concise summary of a patent's core contribution and provide a standardized, comparable text block for analyzing content similarity. The following section outlines my methodology for measuring patent novelty in greater detail.

2.1 Measuring Patent Novelty

The SBERT model developed by Reimers and Gurevych (2019) is a modification of the pre-trained BERT language model (Devlin et al., 2019), optimized for computing semantic similarity. It achieves state-of-the-art performance on the Semantic Textual Similarity Benchmark, outperforming traditional BERT models in capturing sentence-level semantics. To apply this model for measuring patent similarity, I first embed the abstract of each patent in my sample. The purpose of embedding the textual content is to capture the underlying semantics and subject matter of patents in the form of numerical vectors within the embedding space, thereby enabling the measurement of similarity between patents based on their latent representations.

Traditional methods for vectorizing text, such as term frequency and inverse document frequency (TF-IDF), which rely on word frequency or a bag-of-words approach, have been widely used in the literature

(see Hoberg and Phillips, 2016; Cohen, Malloy and Nguyen, 2020; Kelly et al., 2021). However, more recently, the use of language models for vectorizing textual material to measure similarity has been increasingly adopted in the literature. Graeber, Roth and Zimmermann (2024) and Breitung and Müller (2025) use language models from OpenAI to convert textual data into numerical vector embeddings, which are then used to measure similarities. [7]

I adopt an embedding-based approach for textual analysis because methods based on word frequency overlook sentence structure and, as a result, fail to capture contextual meaning. In contrast, embedding models are designed to capture semantic similarity by encoding both the word content and the contextual relationships among words within a sentence. In this regard, models like SBERT can produce similar embeddings for sentences that use different words to express the same underlying concepts. This enables more accurate comparisons of patent texts beyond mere word matching, making embedding models particularly well-suited for analyzing patents, where related innovations are often described using varied terminology.

Among various language models, I use SBERT in particular, as it is specifically designed to measure similarity between sentences.[8] As a result, SBERT has become one of the most widely adopted models for semantic similarity tasks in natural language processing, both in professional and academic settings. In the vectorization process, SBERT maps the textual content of a patent's abstract into a 768-dimensional embedding space. The dimensionality of the embeddings is determined by the SBERT model and cannot be altered. Consequently, each patent in my sample is represented as a 768-dimensional vector in embedding space.

[7] The model used by Graeber, Roth and Zimmermann (2024) is $text - embedding - 3 - large$ from OpenAI, which transforms text into high-dimensional numerical vectors that capture deep semantic relationships between words, sentences, or documents.

[8] SBERT modifies the BERT network to derive semantically meaningful sentence embeddings that can be compared using cosine similarity. Its architecture produces fixed-sized sentence embeddings, ensuring that semantically similar sentences are positioned closely in the embedding space (Reimers and Gurevych, 2019).

Before calculating pairwise cosine similarities, I normalize the embedding vectors of the patents, as shown below

$$\hat{\mathbf{p}}_k = \frac{\vec{\mathbf{p}}_k}{\|\vec{\mathbf{p}}_k\|} \quad , \quad where \quad \vec{\mathbf{p}}_k = [d_{k,1}, d_{k,2}, ..., d_{k,768}] \tag{1}$$

For patent i, I calculate its average pairwise cosine similarity with the past patents as follows

$$\overline{\text{sim}}_{j<i}(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_j) = \frac{1}{\left|\{ j : t_j < t_i \}\right|} \sum_{j: t_j < t_i} \hat{\mathbf{p}}_i \cdot \hat{\mathbf{p}}_j \tag{2}$$

Let $\hat{\mathbf{p}}_j$ denote the normalized embedding vector of any patent j granted within the five years preceding the grant date of patent i. I compute the average cosine similarity between patent i and all prior patents to measure the degree to which its innovation builds upon past technological developments. Patent novelty is then defined as the inverse of this average similarity

$$\text{novelty}_i = \log\left(\frac{1}{\overline{\text{sim}}_{j<i}(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}_j)}\right) \tag{3}$$

This formulation captures the notion that a patent is more novel when it is less similar to prior art. Following Kelly et al. (2021), I demean the novelty scores by grant year before using them in my analysis to account for year-specific fixed effects. This adjustment controls for variation in factors such as the pace of innovation, citation practices, and language use over time. For the remainder of the study, I refer to the demeaned values derived from Eq. 3 as my measure of patent novelty.

[Insert Table 1 near here]

Panel A of Table 1 presents the distribution of patent novelty scores for semiconductor patents, categorized by the assignee's country. U.S. assignees account for the largest share of granted patents and exhibit

the highest mean and median novelty scores, indicating both a high volume of patenting activity and relatively greater innovative distinctiveness. China, while producing only one-tenth the number of patents granted to U.S. assignees, ranks second in mean novelty, consistent with its more recent emergence as a significant contributor to semiconductor innovation. In contrast, assignees from Japan, South Korea, and Taiwan exhibit lower mean novelty scores and reduced dispersion, as evidenced by lower standard deviations, suggesting a greater emphasis on incremental innovations. These cross-country differences highlight the persistent dominance of the U.S. and the growing prominence of China in the domain of semiconductor research and development.

Panel B of Table 1 presents the distribution of patent novelty for U.S. public firms, grouped by market capitalization quintile. Firms in the highest quintile exhibit both the highest mean novelty score and the most pronounced upper-tail value at the 90th percentile. In general, the positive relationship between market capitalization and mean, median, and upper-tail novelty indicates that larger firms contribute disproportionately to innovations characterized by greater novelty. This finding is consistent with the notion that larger firms have greater capacity for research and development (R&D) expenditures, allowing them to invest in a wider range of research activities and undertake riskier projects with higher novelty potential.

2.2 Patent Novelty and Patent Quality

The innovation literature commonly uses forward patent citations as a proxy for patent impact and technological value (see Kogan et al., 2017). In this section, I examine whether my ex ante measure of patent novelty predicts subsequent patent citations. I benchmark my measure against that of Kelly et al. (2021) (hereafter $novelty^{KPST}$), which is based on cosine similarity computed from TF-IDF vectors. This comparison is informative because TF-IDF representations abstract from word order and semantic context,

whereas sentence embeddings capture contextual relationships across words and phrases. Differences in predictive power therefore speak directly to the value of incorporating richer semantic information. I estimate the following cross-sectional regression:

$$\log(1 + fut.cites_i) = \alpha + \beta_{nov}\, novelty_i + \gamma_{assignee} + \delta_{year} + \varepsilon_i$$

where patent novelty is alternatively measured using my proposed measure and $novelty^{KPST}$, and the specification includes assignee and year fixed effects. I obtain patent citation data and $novelty^{KPST}$ scores from the supplemental materials accompanying Kelly et al. (2021) for the semiconductor patents in our sample. Forward citations are measured using information available through the end of 2022.

[Insert Table 2 near here]

Table 2 reports the results. My novelty measure is a statistically significant predictor of future citations across all specifications, including those with both year and assignee fixed effects. By contrast, $novelty^{KPST}$ is not statistically significant in any specification. I note that Kelly et al. (2021) do not use $novelty^{KPST}$ in isolation, but instead combine it with a forward-looking impact measure to identify breakthrough patents ex post. My finding that embedding-based novelty predicts citations suggests that embedding vectors capture dimensions of semantic content that word vector approaches overlook.

2.3 Semiconductor Innovation Index

Having developed my measure of patent novelty, I now construct a semiconductor innovation index by aggregating patent-level novelty over time. I interpret patent novelty as a proxy for the underlying intensity of innovation, consistent with the World Intellectual Property Organization's (WIPO) characterization of innovation as the development of a novel solution to a technical problem. The innovation index is defined

as the share of patents falling within the top decile of the novelty distribution—which I refer to as truly novel patents—relative to total patents granted in each period, following the approach of Kogan et al. (2017).

[Insert Fig. 4 near here]

Figure 4 plots the evolution of semiconductor patenting activity from 1981 to 2023. Each bar represents total patents granted in a given year; the lighter-shaded segment denotes patents in the top decile of my novelty distribution. The figure reveals a persistent upward trend in patenting since the early 1980s, with a pronounced acceleration during the late 1990s—coinciding with the information technology boom—and a renewed surge after 2008, consistent with intensified R&D investment in advanced chip design and manufacturing. The expanding lighter-shaded component indicates that periods of heightened patenting are also characterized by a greater incidence of truly novel innovations, suggesting that expansions in the scale of inventive effort are accompanied by substantive technological progress rather than merely incremental additions to the patent stock.

The value of my innovation index lies in its ability to distinguish meaningful innovation from patenting volume, ex-ante. The literature documents that firms may patent defensively to secure bargaining positions in cross-licensing negotiations or deter infringement litigation (see Hall and Ziedonis, 2001; Ziedonis, 2004), or accumulate patents as signals of technological capability to investors and creditors, rather than as reflections of genuine technological advance (see Abrams, Akcigit and Popadak, 2013; Mann, 2018; Farre-Mensa, Hegde and Ljungqvist, 2020). By conditioning on novelty, my index filters out such activity and isolates the economically meaningful innovation.[9]

[9] Appendix A2 provides a detailed examination showing that the semiconductor innovation index captures economically meaningful technological progress, as measured by subsequent gains in total factor productivity growth across industries.

3. Semiconductor Innovations and Prices of Durables

3.1 Cost of Chip Production

I begin by examining the effect of semiconductor innovation on chip production costs. I hypothesize that, in the short run, this is the primary channel through which semiconductor innovation affects the prices of chip-intensive durables. In the innovation literature, cost-reducing innovation is commonly referred to as *process innovation*, which is distinct from *product innovation*, whose primary objective is to improve product performance or functionality (see Abernathy, 1975; Klepper, 1996; Adner and Levinthal, 2001; Tham et al., 2026). Process innovation introduces new production methods that improve manufacturing efficiency and reduce production costs. I expect the process-innovation channel to dominate in the short run because translating advances in semiconductor technology into pricing of durables through higher performing final goods—that is, product innovation—requires downstream firms to undertake additional R&D, design, and product development efforts around the new more capable chips, a process that may take years.

To demonstrate that semiconductor innovations have such effect on chip production, I use PPI^{semi} monthly series representing the Producer Price Index for Semiconductor Components as a proxy for chip production costs, because it reflects price movements at the producer level and thus provides a natural measure of production-side cost conditions in the semiconductor industry. I estimate the monthly time series below:

$$\Delta PPI_t^{semi} = \alpha + \beta_k^{\text{semi}} \Delta semi_idx_{t-k} + \beta^c C_{t-1} + \varepsilon_t. \tag{4}$$

where the log growth of the seasonally adjusted PPI^{semi} is regressed on the lagged log growth of the

seasonally adjusted semiconductor innovation index. The vector C_{t-1} includes lagged controls for semiconductor industry capacity utilization, to capture supply-side tightness, the producer price index for all commodities, to isolate semiconductor-specific pricing dynamics from broad-based inflationary pressures, total industrial production, to capture aggregate demand pressures, and energy prices and interest rates to control for the effects of input cost and borrowing costs on pricing.

[Insert Table 3 near here]

Table 3 reports the regression results. Across specifications, the innovation index enters with a negative and statistically significant coefficient at the second lag ($\Delta semi_idx_{t-2}$), significant at the 1% level. This finding indicates that major semiconductor innovations reduce chip production costs with a delay of approximately two months. The result is consistent with the view that innovations, especially process innovations, require a short implementation period before their cost-reducing effects appear in producer prices. This short lag is plausible because process innovations directly improve fabrication efficiency as they enable chip fabrication facilities to produce more defect-free chips from a same quantity of inputs, allowing unit costs to decline within a short period.[10]

3.2 Prices of Chip-intensive Durables

Having established that semiconductor innovations reduce chip production costs, I next examine whether these cost reductions pass through to the prices of durables that use chips as production inputs. I focus on two major categories of chip-intensive durables: motor vehicles and electronics/appliances. As a first step, I estimate local projections following Jordà (2005) to identify the horizon over which changes in

[10]Industry reports from ASML reveals that innovations like EUV lithography reduces the number of process steps and defect rates, while TSMC's 7nm node ramped rapidly to over one billion defect-free chips—both mechanisms that can compress per-unit costs relatively quickly (ASML, 2026).

chip prices are transmitted to consumer durable prices. These estimates provide a benchmark for the expected lag between semiconductor innovation shocks and downstream price adjustment in durable goods markets. I stress that this exercise is not intended to isolate only the component of chip price movements attributable to semiconductor innovation. Rather, its purpose is to characterize the transmission dynamics from chip input prices to durable goods prices, regardless of the underlying source of the chip price change.

[Insert Fig. 5 near here]

Figure 5 presents the local projection estimates of the responses of consumer price inflation for motor vehicles and electronics to a shock in PPI^{semi}. The solid lines plot the estimated coefficients at each horizon, while the dashed lines report confidence bands. In the top plot, the dependent variable is the monthly log growth in $CPI^{vehicle}$, which is the *Consumer Price Index for All Urban Consumers: New Vehicles in U.S. City Average*. In the bottom plot, the dependent variable is the monthly log growth in $CPI^{electronics}$, which is the *Consumer Price Index for All Urban Consumers: Computers, Peripherals, and Smart Home Assistants in U.S. City Average*. Both series are obtained from the Federal Reserve Bank of St. Louis's FRED database.

The shock is defined as a one-standard-deviation innovation in the log growth of PPI^{semi} that is orthogonal to the predictors included in the projection. Specifically, prior to estimation, ΔPPI_t^{semi} is purged of predictable variation by regressing it on the full set of controls and retaining the residual, so that the identifying shock is orthogonal to the predictors of the dependent variables. The estimates indicate that semiconductor price movements pass through to electronics prices relatively quickly, with the strongest response occurring at a horizon of roughly three months. For motor vehicles, the response emerges more gradually, with the point estimates suggesting the largest effect at around months six to

eight. In Table 4, Panels A and B I report the results from time series regressions below, where I use $CPI^{vehicle}$ and $CPI^{electronics}$ as the dependent variables:

$$\Delta CPI_t = \alpha + \beta_k^{\text{semi}} \, \Delta\big(PPI_{t-k}^{\text{semi}}\big) + \beta^c \, C_{t-1} + \varepsilon_t. \tag{5}$$

The regression results closely align with the local projection evidence in Figure 5. Consistent with the impulse responses, changes in PPI^{semi} precede movements in $CPI^{vehicle}$ and $CPI^{electronics}$ with lags of approximately seven and three months, respectively. When combined with the findings in Table 3, which show that semiconductor innovations reduce PPI^{semi} with a lag of about two months, these estimates imply that semiconductor innovations lower vehicle prices over an approximately nine-month horizon. Panel A corroborates this mechanism. In addition, once both PPI^{semi} and $semi_idx$ are included, the coefficient on the innovation index becomes insignificant, indicating that the effect operates primarily through reductions in chip input costs.

[Insert Table 4 near here]

In contrast, this cost-based channel appears considerably weaker for $CPI^{electronics}$. In Panel B, the innovation index does not exhibit a robust independent effect, suggesting that semiconductor innovation influences consumer electronics prices less through cost reductions and more through improvements in chip performance and capability. A plausible interpretation is that the automotive sector relies more heavily on legacy chips—older-generation semiconductor components—for which innovation is predominantly directed toward enhancing production efficiency rather than pushing the performance frontier. Consumer electronics, by contrast, are more tightly linked to advancements in chip quality and functionality, implying that their price dynamics reflect a broader set of innovation channels that likely materialize over longer horizons.

Taken together, the evidence in this section traces an economic channel through which semiconductor innovation propagates to the prices of chip-intensive durables, primarily motor vehicles. The next section examines whether these innovation-driven price dynamics affect the timing of household vehicle purchases. Specifically, if consumers anticipate that semiconductor innovations will lower vehicle prices, they may delay purchases, a forward-looking timing response that links innovation cycles to aggregate consumption dynamics.

4. Semiconductor Innovations and Durables Consumption

This section examines how semiconductor innovations affect aggregate consumption through their impact on durable goods demand. Building on the evidence that innovations lower chip costs and reduce durable goods prices, I test whether these changes translate into shifts in consumer expectations and purchasing behavior and timing. Establishing this consumption channel is central to linking semiconductor innovations to asset pricing implications in the subsequent section.

4.1 Consumer Price Expectations for Durables

I examine consumer expectations for durable goods prices using data from the University of Michigan Surveys of Consumers, a long-running, nationally representative monthly survey that tracks household attitudes toward economic conditions, personal finances, and purchasing decisions. The survey is based on a rotating cross-section of U.S. households and includes a core set of approximately 50 questions that remain consistent over time, enabling the construction of stable time-series measures of consumer sentiment and expectations. Importantly, several questions elicit respondents' views on whether it is a good or bad time to purchase major household durables and vehicles. These questions are supplemented by questions that ask respondents to specify the reasoning behind their views, allowing for a more granular interpretation of underlying expectations.

I focus on these "buying conditions" questions as forward-looking indicators of consumer price expectations for durable goods. In particular, unfavorable buying conditions are often associated with expectations of future price declines or anticipated improvements in product quality, both of which incentivize households to delay purchases. I exploit both the responses measuring the share of respondents reporting good or bad buying conditions, and the reasoning responses, to better understand the drivers of these views.

I begin by estimating time-series regressions of the percentage of respondents reporting favorable and unfavorable buying conditions against the semiconductor innovation index. The results are reported in Table 5. Across specifications, increases in the log growth of the innovation index are associated with a decline in the percentage of respondents who view current conditions as favorable and a corresponding increase in the percentage who view them as unfavorable for purchasing large household durables. I obtain similar results for survey questions specific to vehicle purchases. Notably, these effects emerge with a lag of approximately two months, consistent with the timing at which semiconductor innovations begin to affect chip production costs based on the results in Table 3. Taken together, the findings suggest that consumers incorporate information about technological progress into their expectations, revising their assessments of current purchasing conditions in anticipation of lower future prices for durable goods.

[Insert Table 5 near here]

Table 6 provides further insight into the mechanisms underlying consumers' unfavorable assessments of vehicle purchasing conditions by decomposing survey responses based on the stated reasons. Among the different categories, the results indicate that semiconductor innovation primarily operates through price-related expectations. In particular, the two-month lag of the innovation index enters with a positive and statistically significant coefficient in columns (2), implying that increases in semiconductor innovation

are followed by a higher share of respondents citing *high prices* as the reason for viewing current conditions as unfavorable. By contrast, the coefficients on the innovation index are economically small and statistically insignificant for reasons related to uncertainty about the future, affordability constraints, and poor product selection. This pattern suggests that the response documented in Table 5 is not driven by heightened macroeconomic uncertainty or deteriorating household balance sheets, but rather by forward-looking price expectations. Consistent with a wait-and-see mechanism, consumers expect future price declines, which makes current prices seem relatively high and leads them to postpone vehicle purchases.

[Insert Table 6 near here]

I argue that it is unlikely that consumers respond directly to semiconductor innovations themselves, and instead respond to declines in chip costs and improvements in technology, which are reflected in news coverage, retail pricing trends, firm communications, and industry reports. These channels transmit the implications of innovation to households, prompting them to revise their expectations of future durable goods prices downward.

4.2 Demand for Motor Vehicles

I continue my analysis by examining monthly motor vehicle sales data to assess whether changes in consumer price expectations translate into observable shifts in demand. To this end, I rely on both the dollar value of vehicle retail sales and the quantity of vehicles sold to track aggregate demand. Specifically, I consider the deflated series *Retail Sales: New Car Dealers*, as well as unit-based measures including *Total Vehicle Sales* and *Light Weight Vehicle Sales: Autos and Light Trucks*, all obtained from the Federal Reserve Bank of St. Louis's FRED database. Consistent with my earlier findings, I expect a decline in demand following periods of heightened semiconductor innovation, as these innovations lower chip input costs and are subsequently passed through to vehicle prices, inducing consumers to delay purchases.

[Insert Table 7 near here]

Table 7 reports the results from the time-series regressions. Across specifications, the estimates indicate that both the value of new car sales and the total number of vehicles sold decline significantly following semiconductor innovation shocks, with the strongest effects emerging at a two-month lag. Notably, the one-month lag of the innovation index is small and statistically indistinguishable from zero across all specifications, confirming that the demand response concentrates precisely at the two-month horizon and is consistent with the timing at which semiconductor innovations reduce chip production costs. In economic terms, a one standard deviation increase in the innovation index is associated with a reduction in the monthly growth rate of new car sales equivalent to approximately 10% of its standard deviation, and a corresponding reduction of about 12% for the monthly growth rate in total vehicle units sold. The contraction is also evident in light vehicle sales, which comprise the bulk of the U.S. motor vehicle market. These results are consistent with the revision in consumer price expectations documented earlier, together suggesting a wait-and-see mechanism. [11]

In contrast, and consistent with the results in Table 4, semiconductor innovations do not appear to affect the demand for consumer electronics through the cost channel. As shown in Table A3 in the Appendix, retail sales of electronics do not respond significantly to the innovation index at comparable horizons. This further supports the interpretation that the price expectation channel operates primarily in the motor vehicle sector, where semiconductor inputs constitute a critical cost component.

[11] The decline in vehicle sales could also reflect supply-side disruptions in motor vehicle production associated with semiconductor-related developments. To address this concern, I examine whether semiconductor innovations predict changes in retail inventories and manufacturing shipments for motor vehicles in Table A4 in the Appendix. I find no evidence supporting a supply-driven interpretation.

4.3 Aggregate Consumption Movements

The key link between semiconductor innovations and the cross-section of stock returns operates through their effect on short-run consumption dynamics. In this section, I examine whether the decline in demand for motor vehicles is sufficiently large to generate a measurable slowdown in aggregate consumption growth. Figure 6 plots the expenditure shares of two major chip-intensive durables— motor vehicles and electronics—within total durables spending, as well as their shares in overall personal consumption expenditure (PCE). The figure shows that motor vehicles constitute a large and relatively stable share of durable goods expenditures, whereas electronics and appliances account for a substantially smaller fraction. These patterns motivate my analysis in this section, as fluctuations in motor vehicle expenditures are more likely to translate into meaningful variation in aggregate durables consumption and, in turn, overall PCE growth.

[Insert Fig. 6 near here]

I estimate the short-run effect of semiconductor innovations on monthly growth in PCE and its three categories of durables, non-durables, and services. Eq. 6, estimates monthly growth in PCE against lagged growth in the semiconductor innovation index, controlling for standard macroeconomic variables including lagged consumption and income growth, energy prices, financial conditions, and measures of macroeconomic uncertainty.

$$\Delta PCE_t = \alpha + \beta_k^{\text{semi}} \, \Delta(\text{semi_idx}_{t-k}) + \beta^c \, C_{t-1} + \varepsilon_t. \tag{6}$$

Table 8 reports the results. A one standard deviation increase in semiconductor innovation growth is

associated with a decline in monthly PCE growth of approximately 3 percent of its standard deviation. This effect is primarily driven by durable goods consumption, where the lagged coefficient on the innovation index is negative and statistically significant, with the strongest response occurring after two months, consistent with my earlier evidence on the timing of the innovation transmission channel. A one standard deviation increase in semiconductor innovation growth is associated with a decline in the monthly growth of durable goods expenditure of approximately 7.5 percent of its standard deviation. By contrast, the response of nondurable consumption is economically small and statistically insignificant. Despite the smaller expenditure share of durables relative to services, the magnitude of their decline is sufficiently large to drive the statistically significant reduction in overall PCE growth.

[Insert Table 8 near here]

Although the estimated monthly effect of growth in semiconductor innovation on PCE growth is small in magnitude, the persistent nature of innovation activity implies a nontrivial cumulative impact on aggregate consumption. As illustrated in Figure 7, growth in the semiconductor innovation index has exhibited strong persistence over time, with an increasingly steeper slope in the cumulative series since the early 2010s, indicating a sustained acceleration in innovation intensity. This persistence reflects the cumulative nature of technological progress, in which each generation of semiconductor advances lowers the cost and raises the likelihood of subsequent innovations. Repeated innovation shocks, therefore, compound over longer horizons, implying that persistent semiconductor innovation can produce an economically meaningful slowdown in consumption growth in the short term.

[Insert Fig. 7 near here]

In summary, the evidence in this section points to a coherent transmission mechanism through which semiconductor innovations affect aggregate consumption. Innovations lead to a rapid adjustment in con-

sumer price expectations, particularly for motor vehicles, inducing households to delay purchases. This shift in intertemporal demand is reflected in a pronounced contraction in durable goods consumption, which, despite its smaller expenditure share, is sufficiently large to generate a measurable slowdown in overall PCE growth. These findings establish a direct link between technological progress in semiconductors and short-run consumption dynamics. The following section examines how such innovation-driven fluctuations in aggregate consumption translate into asset pricing implications.

5. Semiconductor Innovation Shocks and Stock Returns

The sharp decline in durable goods demand, particularly in motor vehicles, links semiconductor innovation shocks to aggregate consumption risk in the standard Consumption Capital Asset Pricing Model (CCAPM) framework. When semiconductor innovations trigger a slowdown in aggregate consumption growth, they correspond to states of high marginal utility. Assets whose payoffs covary positively with these innovation shocks therefore provide consumption insurance and thus in equilibrium, command lower expected returns. This section tests this prediction by estimating the cross-sectional price of risk associated with exposure to semiconductor innovation shocks, thereby connecting the real effects of technological progress in semiconductors to the pricing of risk in the equity market. My sample comprises monthly stock returns for publicly traded firms from January 1981 to December 2023, matching the coverage of my innovation index. The returns data are obtained from the Center for Research in Security Prices (CRSP).

I exclude the stocks of 320 firms operating in the semiconductor industry from my sample, as they are likely the direct assignees of patents underlying my semiconductor innovation index.[12] The stock returns of such firms are closely tied to the expected cash flows arising from their own patented innovations, reflecting a firm-specific cash-flow channel. Consequently, their estimated factor loadings on the semi-

[12]Specifically, these are firms classified under the Semiconductors & Semiconductor Equipment category according to the 4-digit Global Industry Classification Standard (GICS) code.

conductor innovation index would predominantly capture direct cash-flow effects rather than systematic exposure to innovation-induced consumption growth risk.

I use unexpected shocks to the semiconductor innovation index (henceforth, innovation shocks) to estimate stock-level factor loadings. The estimation of the corresponding price of risk follows a two-stage procedure. In the pre-ranking stage, I estimate firm-level loadings on innovation shocks and market excess returns each month, using stock return data from the preceding 12 months. Firms are then sorted into 100 portfolios based on these estimated loadings. In the first stage, I estimate portfolio-level factor loadings on innovation shocks, the five factors of Fama and French (2015), and the momentum factor of Carhart (1997), using returns over the previous 30 months. In the second stage, I estimate the cross-sectional variation in average stock returns explained by these first-stage factor loadings. The resulting coefficient represents the estimated price of risk associated with exposure to semiconductor innovation shocks.

5.1 Pre-ranking Factor Loadings

Each month, I estimate the following time-series regression at the stock-level using a rolling window of the past 12 monthly returns, excluding the current month. This specification yields two factor loadings for each stock-month observation, available ex-ante at the beginning of the month:

$$r_{i,t} = \alpha_{i,t} + \beta_{i,t}^{\text{semi}} \, \Delta(\text{semi_idx}_\text{t}) + \beta_{i,t}^{\text{mkt}} \, MKT_t + \varepsilon_{i,t} \tag{7}$$

where $r_{i,t}$ denotes the excess return of stock i over the risk-free rate. The variable semi_idx_shock_t represents the unexpected shocks to my innovation index at month t. The value of the shock series at month t is the last residual collected from an AR(1) model applied to $semi_idx$ series up to the current

month. I extract $\beta_{i,t}^{\text{semi}}$ and $\beta_{i,t}^{\text{mkt}}$ as factor loadings on innovation shocks and the market risk premium, respectively. Using these estimated loadings, I form 100 equally weighted portfolios via double-sorting on $\beta_{i,t}^{\text{semi}}$ and $\beta_{i,t}^{\text{mkt}}$. Portfolios are rebalanced annually based on the factor loadings available at the end of June.

In Table 10, I report the median pre-ranking loadings on innovation shocks by quantiles constructed based on size, market beta, and book-to-market. The results indicate that exposure to innovation shocks is most strongly related to market betas, with higher beta stocks exhibiting larger loadings. I interpret the stronger loadings of higher–market-beta stocks on innovation shocks through the durability channel. Since demand for durable goods is more cyclical than for nondurables, the cash flows and expected returns of firms linked to durable goods production are more sensitive to aggregate market conditions. Moreover, because semiconductor advances are embodied in durable goods, such innovations directly affect the future cash flows of these firms. Consequently, a stronger association between market beta and exposure to innovation shocks naturally arises for firms whose performance is closely tied to cyclical, chip-intensive durables demand. In contrast, the relationship between innovation loadings and firm size or book-to-market ratio is less systematic, showing no consistent monotonic pattern across quintiles. This implies that the price of risk associated with semiconductor innovations is not linked to the distinction between value and growth assets as suggested by the literature (see Kung and Schmid, 2015; Kogan, Papanikolaou and Stoffman, 2020).

[Insert Table 9 near here]

5.2 Two-Stage Regressions

To obtain portfolio-level factor loadings, I estimate monthly exposures to innovation shocks, the five factors of Fama and French (2015), and the momentum factor of Carhart (1997). Specifically, for each

portfolio, I run the following time-series regression each month using a rolling window of the past 30 monthly returns, excluding the current month:

$$\begin{aligned} r_{p,t} = \alpha_{p,t} + \beta_{i,t}^{\text{semi}}\,\Delta(\text{semi_idx}_{\text{t}}) + \beta_{p,t}^{\text{mkt}}\,MKT_t + \beta_{p,t}^{\text{smb}}\,SMB_t \\ + \beta_{p,t}^{\text{hml}}\,HML_t + \beta_{p,t}^{\text{rmw}}\,RMW_t + \beta_{p,t}^{\text{cma}}\,CMA_t + \beta_{p,t}^{\text{umd}}\,UMD_t + \varepsilon_{p,t} \end{aligned} \tag{8}$$

where $r_{p,t}$ denotes the excess return of portfolio p over the risk-free rate. I extract $\beta_{p,t}^{\text{semi}}$ as the portfolio's loading on the innovation shocks available ex-ante at the beginning of month t.

The following is the regression specification for the second-stage estimations:

$$\begin{aligned} r_{p,t} = \alpha_t + \boldsymbol{\pi_t^{\text{semi}}}.\boldsymbol{\beta_{p,t}^{\text{semi}}} + \pi_t^{\text{mkt}}.\beta_{p,t}^{\text{mkt}} + \pi_t^{\text{smb}}.\beta_{p,t}^{\text{smb}} + \pi_t^{\text{hml}}.\beta_{p,t}^{\text{hml}} + \\ + \pi_t^{\text{rmw}}.\beta_{p,t}^{\text{rmw}} + \pi_t^{\text{cma}}.\beta_{p,t}^{\text{cma}} + \pi_t^{\text{umd}}.\beta_{p,t}^{\text{umd}} + \varepsilon_{p,t} \end{aligned} \tag{9}$$

r_p is the monthly excess return over the risk-free rate for portfolio p. π^{semi} represents the price of innovation risk and is estimated separately for each month. This procedure yields a monthly time series of π^{semi} estimates from July 1987 to December 2023. To assess the statistical significance of the estimated price of risk, I compute Newey-West t-statistics based on the time-series average of π^{semi} and its corresponding standard error. Appendix 6 outlines the methodology used to estimate Newey-West t-statistics for the time-series averages of the prices of risk.

In Table 11, Columns (1) to (4) report specifications where both the first- and second-stage portfolios are formed using stocks from all industries. Columns (5) to (8)instead restrict the analysis to firms in chip-dependent industries, Consumer Durables, Manufacturing, Business Equipment (Technology), and Telecom, based on the Fama–French twelve industry classification. Focusing on these industries improves

measurement precision, as pre-ranking estimates of stock-level loadings on semiconductor innovation shocks are less noisy for firms whose operations and cash flows are more directly linked to semiconductor technologies. Excluding firms from sectors with limited semiconductor exposure in forming the 100 first-stage portfolios mitigates estimation noise for β_p^{semi}. This approach enhances the cross-sectional signal and identification of the price of risk associated with semiconductor innovation shocks.

The results in Columns (1) to (4) show that the cross-sectional explanatory power of β^{semi} overlaps with that of other factors. When considered alone, the estimated price of risk is insignificant. Once I account for these effects by including stocks' loadings on the standard risk factors, the loadings on innovation shocks exhibit significant explanatory power in the cross-section of stock returns with the expected sign. For the chip-dependent sample in Columns (5) to (8), the estimated price of risk is negative and statistically significant before and after adding controls. A potential concern here is that restricting to chip-dependent industries undermines the notion that a price of risk reflects marketwide, systematic exposure. However, if semiconductor innovation risk were not truly priced, investors could span or hedge it using non-chip-dependent industries, and no significant price of risk would emerge, even within chip-intensive sectors. The fact that I find a significant price of risk in the chip-dependent subsample, therefore, points to the presence of a systematic component, and the failure to detect it in the sample of stocks from all industries can be attributed to attenuation from noisier pre-ranking beta estimates that dilute the cross-sectional signal.

[Insert Table 10 near here]

In Fig. 8, I present the annual averages of the estimated monthly price of risk for the sample that uses stocks from all industries. Dark bars denote averages computed using only those months in which the price of risk is statistically significant, whereas lighter bars include all months. The estimated price of risk is significantly negative in 22 out of 36 years. The persistence of negative values over time suggests

that the results are not driven by a few episodic observations but rather reflect a systematic and recurring pricing pattern.

[Insert Fig. 8 near here]

The results from the two-stage cross-sectional regressions confirm the prediction that stocks with higher exposures to semiconductor innovation shocks earn lower expected returns, consistent with these shocks coinciding with states of lower consumption growth.Taken together, the evidence indicates that semiconductor innovation shocks represent a systematically priced source of consumption growth risk in the cross-section of equity returns.

6. Conclusion

This paper provides new evidence on the short-run economic effects and asset pricing implications of semiconductor innovation. I construct a semiconductor innovation index from the textual novelty of patents in semiconductor devices and use it to track major innovation episodes over time. I show that semiconductor innovation lowers the cost of chip production, and that these cost reductions pass through to the prices of major durable goods—particularly motor vehicles—within a few months.

I then examine how households respond to this channel. Combining survey evidence on expected future durables prices with aggregate retail sales for new vehicles, I find that major semiconductor innovation episodes are followed by downward revisions in price expectations and a decline in new vehicle demand. These patterns are consistent with a wait-and-see mechanism in which households postpone purchases of big-ticket durables in anticipation of lower future prices. The resulting shift in durables expenditure is large enough to contribute to a temporary slowdown in monthly aggregate consumption growth.

This dynamic links semiconductor innovations to aggregate consumption movements, a key determinant of asset returns under the CCAPM. Consistent with this, I find that stocks with higher exposure to

semiconductor innovation shocks earn lower expected returns, as their payoffs is higher in those states with depressed consumption and higher marginal utility of consumption following innovation shocks. Using a Fama–MacBeth two–stage regression framework, I estimate a negative and statistically significant price of risk associated with semiconductor innovations that persists across decades. Taken together, these results contribute to the literature on the economic impact of innovation. While prior work emphasizes on the role of innovation as the main driver of long-run productivity growth, my findings highlight its role as a source of priced short-run consumption risk in equity markets.

Tables

Table 1: Summary Statistics of Patent Novelty for Semiconductor Patents

Panel A: By Country of Assignee

Country	Num. Obs.	Mean	Median	SD	P10	P25	P75	P90
US	165609	0.0407	0.00319	0.228	-0.186	-0.102	0.135	0.292
China	16054	0.0183	0.0000145	0.200	-0.221	-0.110	0.127	0.263
Japan	131511	-0.0234	-0.0498	0.191	-0.246	-0.168	0.0887	0.231
S. Korea	43935	-0.0532	-0.0715	0.177	-0.266	-0.185	0.0509	0.173
Taiwan	54473	-0.0469	-0.0694	0.195	-0.274	-0.191	0.0580	0.194
other	59186	0.0178	-0.0119	0.218	-0.222	-0.130	0.127	0.277

Panel B: By Market Cap. of Assignee

(Only for US Assignees that are Public Firms)

MCap. Quintile	Mean	Median	SD	P10	P25	P75	P90
1	0.000661	-0.0349	0.206	-0.209	-0.132	0.0853	0.240
2	-0.00157	-0.0387	0.204	-0.209	-0.134	0.0839	0.236
3	-0.0175	-0.0468	0.184	-0.214	-0.139	0.0659	0.197
4	-0.0171	-0.0457	0.190	-0.216	-0.142	0.0666	0.199
5	0.0211	-0.0101	0.205	-0.186	-0.110	0.110	0.253

Panel A reports the distribution of patent novelty for granted patents by the country of the patent assignee. The sample contains 467800 semiconductor-device patents (CPC 4-digit) from 1981 to 2023. Panel B reports the distribution of patent novelty for the patents granted to the US public firms from 1981 to 2022. This sample includes 133607 patents across 865 firms, where patents are matched with the owner firm based on the patent-firm matching data from the replication materials accompanying Kogan et al. (2017). Market cap quintiles are constructed as of December of the previous calendar year.

Table 2: Patent Citations and Patent Novelty

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$novelty$	0.037***		0.044***	0.010*		0.022***	0.011*
	(6.581)		(6.168)	(2.126)		(3.996)	(2.310)
$novelty^{KPST}$		-0.011.	-0.024**		-0.042***	-0.048***	-0.053***
		(-1.904)	(-3.259)		(-6.931)	(-7.269)	(-8.509)
$Num.\,obs.$	133607	133607	133607	133389	133389	133389	131615
$Adj.\,R^2$	0.264	0.264	0.265	0.117	0.118	0.118	0.343
$Year\,FE$	Y	Y	Y	-	-	-	-
$Assignee\,FE$	-	-	-	Y	Y	Y	-
$Year \times Assignee\,FE$	-	-	-	-	-	-	Y

In the table above, I report the results from the cross-sectional regression below on the sample of granted patents belonging to the manufacturing of semiconductor devices category, based on the 4-digit Cooperative Patent Classification (CPC) code:

$$log(1 + fut.cites_i) = \alpha + \beta_{nov}\, novelty_i + \gamma_{assignee} + \delta_{year} + \varepsilon_i$$

where $fut.cites_i$ is future citations for patent i, available ex-post at the end of 2022. I obtain patent citations from the replication materials accompanying Kelly et al. (2021). The main independent variable, $novelty_i$, captures the novelty of patent i relative to all patents granted in the prior five years, following my proposed methodology to measure patent similarity. For comparison, I include patent novelty and importance constructed by Kelly et al. (2021) (KPST). KPST's novelty measure uses a word vector approach to measure a patent's similarity to prior patents over the past five years. The sample includes patents granted from 1981 to 2022 across 875 assignee organizations. All independent variables are standardized. Standard errors are clustered at the year level for all regressions.

Table 3: Semiconductor Innovation and Producer Price Index for Semiconductors

	Dependent variable: ΔPPI_t^{semi}				
	(1)	(2)	(3)	(4)	(5)
Constant	−0.001 (−0.583)	−0.001 (−0.577)	−0.001 (−0.578)	−0.001 (−0.576)	−0.001 (−0.476)
$\Delta(\text{semi_idx}_{t-1})$	0.001* (1.825)				
$\Delta(\text{semi_idx}_{t-2})$		**−0.001**** (−2.230)	**−0.001**** (−2.231)	**−0.001**** (−2.235)	**−0.001**** (−2.170)
$\Delta(\text{total_semi_patent}_{t-2})$			0.0002 (0.454)		
$\Delta(\text{total_patent}_{t-2})$				0.0002 (0.379)	
$\Delta(\text{kpst_inov_idx}_{t-2})$					−0.001 (−0.351)
ΔPPI_{t-1}^{semi}	0.192*** (4.551)	0.193*** (4.499)	0.194*** (4.526)	0.194*** (4.553)	0.194*** (4.493)
$\Delta Cap.Util_{t-1}^{semi}$	0.014 (1.460)	0.014 (1.461)	0.014 (1.458)	0.014 (1.455)	0.015 (1.536)
$\Delta PPI_{t-1}^{econ.}$	0.021 (0.899)	0.020 (0.882)	0.020 (0.881)	0.020 (0.886)	0.023 (0.985)
$\Delta Ind.Prod_{t-1}^{econ.}$	−0.032** (−2.049)	−0.031** (−2.066)	−0.031** (−2.063)	−0.031** (−2.057)	−0.030** (−2.058)
$FF\ Rate_{t-1}$	0.0001 (0.742)	0.0001 (0.733)	0.0001 (0.730)	0.0001 (0.730)	0.0001 (0.596)
$log(WTI)_{t-1}$	0.00005 (0.065)	0.00004 (0.061)	0.00004 (0.061)	0.00004 (0.059)	−0.00003 (−0.040)
Observations	476	476	476	476	461
Adjusted R^2	0.033	0.036	0.034	0.034	0.033

The table reports estimates from monthly time-series regressions of the form

$$\Delta PPI_t^{semi} = \alpha + \beta_k^{\text{semi}}\,\Delta(\text{semi_idx}_{t-k}) + \beta^c C_{t-1} + \varepsilon_t.$$

The dependent variable, ΔPPI_t^{semi}, is the monthly log growth of the Producer Price Index for Semiconductor and Other Electronic Component Manufacturing, obtained from the Federal Reserve Bank of St. Louis FRED database. The main explanatory variable, $\Delta(\text{semi_idx}_{t-k})$, is the seasonally adjusted log growth of the semiconductor innovation index lagged by k months, with $k \in \{1, 2\}$. For comparison, additional specifications include the second lag of the seasonally adjusted log growth in the number of semiconductor patents, $\Delta(\text{total_semi_patent}_{t-2})$; the log growth in the total number of patents across all technology classes, $\Delta(\text{total_patent})$; and the log growth in the aggregate innovation index, $\Delta(\text{kpst_inov_idx})$, constructed from patents across all technology classes using the novelty measure of Kelly et al. (2021). The control vector C_{t-1} includes the lagged dependent variable, lagged log growth in semiconductor-industry capacity utilization, lagged log growth in the Producer Price Index for All Commodities, lagged log growth in the Industrial Production: Total Index, the lagged Federal Funds Effective Rate, and the lagged log of the West Texas Intermediate (WTI) spot price. All control variables are obtained from FRED. Newey–West t-statistics are reported in parentheses. The sample consists of 468 monthly observations from January 1985 through October 2024.

Table 4: Semiconductor Producer Price Index and Consumer Price Index for New Vehicles and Electronics

Panel A: CPI for New Vehicles

	Dependent variable: $\Delta CPI_t^{vehicle}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.001***	0.001***	0.001***	0.001***	0.001***	0.009***
	(5.257)	(4.876)	(4.378)	(3.926)	(4.179)	(2.782)
ΔPPI_{t-7}^{semi}	**0.132*****				**0.153*****	**0.087****
	(3.015)				(3.381)	(2.408)
ΔPPI_{t-8}^{semi}		0.074**				
		(1.999)				
$\Delta(\text{semi_idx}_{t-9})$			**−0.001****		−0.0002	−0.001
			(−2.120)		(−0.463)	(−1.567)
$\Delta(\text{semi_idx}_{t-10})$				0.0001		
				(0.394)		
Observations	481	480	506	505	461	382
Adjusted R^2	0.025	0.006	0.001	−0.002	0.031	0.390
Controls	No	No	No	No	No	Yes

Panel B: CPI for Electronics

	Dependent variable: $\Delta CPI_t^{electronics}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	−0.006***	−0.006***	−0.006***	−0.006***	−0.006***	−0.006
	(−6.432)	(−7.380)	(−5.150)	(−7.388)	(−5.227)	(−0.399)
ΔPPI_{t-3}^{semi}	0.417***				0.383***	0.368**
	(4.545)				(5.464)	(2.554)
ΔPPI_{t-4}^{semi}		0.218*				
		(1.673)				
$\Delta(\text{semi_idx}_{t-5})$			−0.007*		−0.007*	−0.005
			(−1.926)		(−1.879)	(−0.965)
$\Delta(\text{semi_idx}_{t-6})$				0.006		
				(1.438)		
Observations	247	247	227	227	227	226
Adjusted R^2	0.032	0.006	0.002	0.001	0.029	0.066
Controls	No	No	No	No	No	Yes

Panel A reports estimates from Eq. 5, where the dependent variable, $\Delta CPI_t^{vehicle}$, is the monthly log growth of Consumer Price Index for All Urban Consumers: New Vehicles in the U.S. City Average. $\Delta(\text{PPI}^{\text{semi}}{}_{t-k})$ is the log growth of the semiconductor producer price index at lags $k \in \{7, 8\}$. *Panel B* reports estimates from Eq. 5, where the dependent variable, $\Delta CPI_t^{electronics}$, is the monthly log growth of Consumer Price Index for All Urban Consumers: Consumer Electronics. $\Delta(\text{PPI}^{\text{semi}}{}_{t-k})$ is the log growth of the semiconductor producer price index at lags $k \in \{3, 4\}$. In both panels, $\Delta(\text{semi_idx})$ is included at two lags preceding $\Delta(\text{PPI}^{\text{semi}})$ in the specification. The control vector C_{t-1}, included in column (6) of both panels, comprises the lag of the dependent variable; lagged log growth of capacity utilization in the semiconductor industry; lagged log growth of Producer Price Index by Commodity: All Commodities; lagged log growth of Total Vehicle Sales; lagged log growth of Retail Inventories/Sales Ratio: Motor Vehicle and Parts Dealers; lagged log growth of Consumer Price Index for All Urban Consumers: All Items in U.S. City Average; lagged log growth of *Consumer Price Index for All Urban Consumers: Gasoline (All Types) in U.S. City Average*; the lagged *CBOE Volatility Index* (VIX); the lagged *Economic Policy Uncertainty* (EPU) index; the lagged *Federal Funds Effective Rate*; and the lagged *University of Michigan's Consumer Sentiment* index. All controls are obtained from the Federal Reserve Bank of St. Louis's FRED database. Parentheses report Newey–West t-statistics.

Table 5: Semiconductor Innovation and Consumer Surveys on Buying Conditions for Durables and Vehicles

	Dependent Variables:							
	$(\%)GoodTime^{Durables}$		$(\%)BadTime^{Durables}$		$(\%)GoodTime^{Vehicles}$		$(\%)BadTime^{Vehicles}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	13.792***	13.764***	−9.008*	−9.144*	6.620**	6.465**	−4.343*	−4.257*
	(3.384)	(3.401)	(−1.815)	(−1.873)	(2.166)	(2.135)	(−1.697)	(−1.689)
$\Delta(\text{semi_idx}_{t-1})$	1.556*		−1.291		0.463		−0.940	
	(1.930)		(−1.545)		(0.561)		(−1.111)	
$\Delta(\text{semi_idx}_{t-2})$		**−1.945****		**2.258****		**−1.853****		**2.146****
		(−2.624)		(2.978)		(−2.002)		(2.406)
$\log(\text{WTI})_{t-1}$	−1.462***	−1.465***	1.798***	1.790***	−0.930**	−0.913**	1.045***	1.030***
	(−3.802)	(−3.867)	(3.882)	(3.911)	(−2.425)	(−2.400)	(2.941)	(2.921)
$FF\,\text{Rate}_{t-1}$	−0.146**	−0.146**	0.144**	0.145**	−0.061	−0.059	0.050	0.049
	(−2.011)	(−2.072)	(2.094)	(2.157)	(−0.506)	(−0.503)	(0.568)	(0.537)
VIX_{t-1}	−0.076	−0.076	0.065	0.066	0.017	0.016	−0.002	−0.002
	(−1.400)	(−1.413)	(1.253)	(1.272)	(0.735)	(0.717)	(−0.086)	(−0.060)
EPU_{t-1}	−0.008**	−0.008*	0.008**	0.008**	−0.001	−0.001	0.001	0.001
	(−1.967)	(−1.941)	(2.019)	(1.980)	(−0.310)	(−0.235)	(0.322)	(0.267)
$\text{Cons. Sentiment}_{t-1}$	0.007	0.005	0.029	0.030	−0.015	−0.015	0.022	0.022
	(0.177)	(0.144)	(0.743)	(0.791)	(−0.732)	(−0.710)	(1.089)	(1.062)
Observations	407	407	407	407	407	407	407	407
Adjusted R^2	0.908	0.908	0.934	0.935	0.901	0.902	0.930	0.931

This table reports estimates from monthly time-series regressions in which the dependent variables are survey-based monthly percentages on whether it is a good or bad time to buy large household durables and vehicles. Columns (1) and (2) report the percentage of consumers who believe it is a good time to buy large household durables, and Columns (3) and (4) report the percentage who believe it is a bad time to buy large household durables. Columns (5) and (6) report the percentage of consumers who believe it is a good time to buy vehicles, and Columns (7) and (8) report the percentage who believe it is a bad time to buy vehicles. The survey data are from the University of Michigan Surveys of Consumers. All survey percentages are deseasonalized. $\Delta(\text{semi_idx}_{t-k})$ denotes the seasonally adjusted log growth in the semiconductor innovation index in month $t-k$, where $k \in 1, 2$. All specifications include the lag of the dependent variable as control. All other control variables are obtained from the Federal Reserve Bank of St. Louis (FRED). Newey–West t-statistics are reported in parentheses. The sample includes 407 monthly observations from 1981 through 2023.

Table 6: Semiconductor Innovation and Consumer Surveys on Reasons for Opinions about Buying Conditions for Vehicles

	Dependent variables:							
	$(\%)BadTime^{High\ Prices}$		$(\%)BadTime^{Uncertain\ Future}$		$(\%)BadTime^{Cannot\ Afford}$		$(\%)BadTime^{Poor\ Selection}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	3.070	3.159	−0.609	−0.618	0.364	0.368	−1.437	−1.410
	(1.384)	(1.434)	(−0.262)	(−0.267)	(0.197)	(0.201)	(−1.177)	(−1.153)
$\Delta(\text{semi_idx}_{t-1})$	−0.367		−0.102		−0.259		−0.048	
	(−0.633)		(−0.299)		(−0.666)		(−0.176)	
$\Delta(\text{semi_idx}_{t-2})$		**1.292****		0.008		0.388		0.313
		(2.144)		(0.020)		(0.862)		(1.275)
$\log(\text{WTI})_{t-1}$	−0.206	−0.217	0.417	0.418	0.917***	0.915***	1.057***	1.054***
	(−0.929)	(−0.988)	(1.590)	(1.605)	(3.912)	(3.936)	(4.114)	(4.109)
$FF\,\text{Rate}_{t-1}$	−0.081	−0.082	−0.052	−0.052	−0.008	−0.008	0.061	0.061
	(−1.228)	(−1.228)	(−0.831)	(−0.829)	(−0.226)	(−0.227)	(1.627)	(1.618)
VIX_{t-1}	−0.039***	−0.038**	0.029	0.029	0.022	0.022	−0.003	−0.003
	(−2.633)	(−2.571)	(1.384)	(1.378)	(1.335)	(1.369)	(−0.272)	(−0.256)
EPU_{t-1}	−0.001	−0.002	0.012***	0.012***	0.005*	0.005*	−0.002	−0.003
	(−0.567)	(−0.641)	(4.211)	(4.241)	(1.651)	(1.696)	(−1.389)	(−1.400)
$\text{Cons. Sentiment}_{t-1}$	−0.011	−0.011	−0.008	−0.008	−0.026**	−0.026**	−0.006	−0.006
	(−0.755)	(−0.802)	(−0.547)	(−0.546)	(−2.197)	(−2.192)	(−0.738)	(−0.756)
Observations	407	407	407	407	407	407	407	407
Adjusted R^2	0.949	0.950	0.773	0.773	0.813	0.813	0.629	0.630

This table reports estimates from monthly time-series regressions in which the dependent variables are survey-based monthly percentages capturing reasons consumers cite for viewing it as a bad time to buy vehicles. Columns (1) and (2) report the percentage of consumers who cite high prices as the reason. Columns (3) and (4) report the percentage who cite uncertainty about future conditions. Columns (5) and (6) report the percentage who cite affordability constraints and personal-finance reasons. Columns (7) and (8) report the percentage who cite poor selection of vehicles as the reason. The survey data are from the University of Michigan Surveys of Consumers. All survey percentages are deseasonalized. $\Delta(\text{semi_idx}_{t-k})$ denotes the seasonally adjusted log growth in the semiconductor innovation index in month $t-k$, where $k \in 1, 2$. All specifications include the lag of the dependent variable as control. All other control variables are obtained from the Federal Reserve Bank of St. Louis (FRED). Newey–West t-statistics are reported in parentheses. The sample includes 407 monthly observations from 1981 through 2023.

Table 7: Semiconductor Innovation and Vehicle Sales

	Dependent variable:					
	Δ*Car Sales*$_t$(\$)		Δ*Total Vehicle*$_t$(#)		Δ*Light Vehicle*$_t$(#)	
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.039	0.038	0.100	0.097	0.102	0.099
	(0.933)	(0.928)	(1.493)	(1.452)	(1.468)	(1.428)
$\Delta(\text{semi_idx}_{t-1})$	0.011		0.014		0.013	
	(1.283)		(1.395)		(1.333)	
$\Delta(\text{semi_idx}_{t-2})$		**−0.014*****		**−0.032*****		**−0.032*****
		(−2.876)		(−4.480)		(−4.556)
ΔCar Sales$_{t-1}$	−0.138**	−0.137**				
	(−2.121)	(−2.124)				
ΔTotal Vehicle$_{t-1}$			−0.195**	−0.190**		
			(−2.238)	(−2.197)		
Δ Light Vehicle$_{t-1}$					−0.198**	−0.192**
					(−2.326)	(−2.286)
ΔDisp. Income$_{t-1}$	0.243	0.245	0.268*	0.269*	0.286*	0.286*
	(1.581)	(1.593)	(1.823)	(1.780)	(1.873)	(1.834)
Δ Gas. CPI$_{t-1}$	−0.042	−0.044	0.021	0.017	0.018	0.015
	(−0.574)	(−0.608)	(0.232)	(0.189)	(0.198)	(0.158)
Auto Loan Rate$_{t-1}$	0.00004	0.0001	−0.001	−0.001	−0.001	−0.001
	(0.020)	(0.028)	(−0.522)	(−0.496)	(−0.558)	(−0.531)
log(WTI)$_{t-1}$	−0.012**	−0.012**	−0.017**	−0.017**	−0.017**	−0.017*
	(−2.013)	(−2.004)	(−2.005)	(−1.967)	(−1.971)	(−1.932)
FF Rate$_{t-1}$	0.0002	0.0001	−0.0002	−0.0002	−0.00002	−0.00003
	(0.097)	(0.084)	(−0.064)	(−0.069)	(−0.010)	(−0.015)
VIX_{t-1}	−0.001	−0.001	−0.001*	−0.001*	−0.001*	−0.001*
	(−1.522)	(−1.516)	(−1.702)	(−1.722)	(−1.680)	(−1.699)
EPU_{t-1}	0.0002***	0.0002***	0.0001**	0.0001***	0.0001**	0.0001***
	(3.381)	(3.373)	(2.462)	(2.579)	(2.477)	(2.599)
Cons. Sentiment$_{t-1}$	−0.0001	−0.0001	−0.0004	−0.0004	−0.0004	−0.0004
	(−0.245)	(−0.234)	(−0.927)	(−0.890)	(−0.926)	(−0.890)
Observations	382	382	407	407	407	407
Adjusted R^2	0.067	0.068	0.053	0.063	0.055	0.064

This table reports estimates from monthly time–series regressions, where dependent variables are ΔCar Sales$_t$, ΔTotal Vehicle$_t$, and ΔLight Vehicle$_t$ are monthly log growth of the seasonally adjusted and deflated *Retail Sales: New Car Dealers*, *Total Vehicle Sales* (units), and *Light Weight Vehicle Sales: Autos and Light Trucks* (units) respectively, obtained from the Federal Reserve Bank of St. Louis's FRED database. $\Delta(\text{semi_idx}_{t-k})$ is the seasonally adjusted log growth of the semiconductor innovation index at lag k, ($k \in \{1, 2\}$). The control vector includes, as applicable to each specification, the lag of the dependent variable; the lagged log growth of *Real Disposable Personal Income*; the lagged *University of Michigan's Consumer Sentiment index*; the lagged log growth of *Consumer Price Index for All Urban Consumers: Gasoline (All Types) in U.S. City Average*; the lagged *CBOE Volatility Index (VIX)*; the lagged *Economic Policy Uncertainty (EPU)* index; the lagged *Federal Funds Effective Rate*. All control series are sourced from FRED. Newey–West heteroskedasticity- and autocorrelation-consistent t-statistics are reported in parentheses. The sample comprises 382 monthly observations from 1992 to 2023 for new car retail sales, and 407 monthly observations from 1990 to 2023 for vehicle sales in units.

Table 8: Semiconductor Innovation and Monthly Personal Consumption Expenditure For All Goods, Durable Goods, Non-Durable Goods, and Services

	Dependent variable:							
	ΔPCE_t^{all}		$\Delta PCE_t^{durables}$		$\Delta PCE_t^{non-durables}$		$\Delta PCE_t^{services}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.010***	0.010***	0.035***	0.035***	0.013***	0.013***	0.006***	0.006
	(8.497)	(8.345)	(5.083)	(5.097)	(4.830)	(4.828)	(3.642)	(0.842)
$\Delta(\text{semi_idx}_{t-1})$	−0.001		0.008		0.001		0.0002	
	(−0.493)		(1.341)		(0.498)		(0.213)	
$\Delta(\text{semi_idx}_{t-2})$		**−0.001****		**−0.008*****		−0.002		−0.001
		(−2.042)		(−3.027)		(−1.478)		(−0.780)
ΔPCE_{t-1}^{all}	0.068	0.067						
	(1.198)	(1.193)						
$\Delta PCE_{t-1}^{durables}$			0.007	0.009	0.066	0.067	0.068	0.068
			(0.034)	(0.048)	(1.002)	(1.026)	(1.245)	(1.270)
$\Delta PCE_{t-1}^{non-durables}$			−0.831	−0.825	−0.700***	−0.701***	−0.398	−0.398
			(−1.523)	(−1.513)	(−3.714)	(−3.749)	(−1.745)	(−1.815)
$\Delta PCE_{t-1}^{services}$			0.131	0.122	0.461***	0.459***	0.431***	0.429***
			(0.695)	(0.659)	(5.948)	(5.927)	(3.874)	(3.979)
$\Delta Disp.\,Income_{t-1}$	0.044**	0.043**	0.127**	0.128**	0.053*	0.053*	0.035**	0.035*
	(2.278)	(2.274)	(2.455)	(2.399)	(1.829)	(1.813)	(2.529)	(1.793)
$log(WTI_{t-1})$	−0.002***	−0.002***	−0.010***	−0.010***	−0.004***	−0.004***	−0.002***	−0.002
	(−7.505)	(−7.409)	(−4.675)	(−4.684)	(−3.638)	(−3.636)	(−2.724)	(−1.404)
$FF\,Rate_{t-1}$	−0.00005	−0.00005	−0.001	−0.001	−0.0004**	−0.0004**	0.00001	0.00001
	(−0.328)	(−0.314)	(−1.077)	(−1.085)	(−2.211)	(−2.217)	(0.103)	(0.073)
VIX_{t-1}	−0.0003	−0.0003	−0.0005**	−0.0005**	−0.0001	−0.0001	−0.0002**	−0.0002**
	(−1.461)	(−1.452)	(−2.319)	(−2.329)	(−1.577)	(−1.593)	(−2.265)	(−2.033)
EPU_{t-1}	0.00004*	0.00004*	0.0001**	0.0001**	0.00002	0.00002	0.00003*	0.00003**
	(1.839)	(1.834)	(2.175)	(2.181)	(1.561)	(1.572)	(1.917)	(2.270)
Observations	407	407	407	407	407	407	407	407
Adjusted R^2	0.057	0.057	0.150	0.150	0.386	0.386	0.283	0.284

The table reports estimates from the following monthly time series regression:

$$\Delta PCE_t = \alpha + \beta_k^{\text{semi}}\,\Delta(\text{semi_idx}_{t-k}) + \beta^c\,C_{t-1} + \varepsilon_t.$$

ΔPCE_t denotes the monthly log growth of real Personal Consumption Expenditure (PCE) for all goods, durable goods, non-durable goods, and services category, each deflated using the consumer price index for their category, all seasonally adjusted and obtained from the Federal Reserve Bank of St. Louis's FRED database. $\Delta(\text{semi_idx}_{t-k})$ is the seasonally adjusted log growth of the semiconductor innovation index in month $t-k$ ($k \in \{1,2\}$). The control vector C_{t-1} includes the lag of the dependent variables, lagged log growth of real Disposable Personal Income, lagged CBOE Volatility Index (VIX) index, lagged Economic Policy Uncertainty (EPU) index, lagged Federal Funds Effective Rate, and lagged log of West Texas Intermediate (WTI) spot price. All the controls are obtained from the Federal Reserve Bank of St. Louis's FRED database. Reported in parentheses are Newey–West t-statistics. The sample covers 407 monthly observations from 1991 through 2023.

Table 9: Stock-level Exposure to Chip Prices and Semiconductor Innovations

	Dependent variable: $r_{i,t}$				
	Q1	Q2	Q3	Q4	Q5
$\Delta(\text{semi_idx}_t)$	$-$**0.0080****	-0.0038	0.0002	-0.0019	0.0024
	(-2.25)	(-1.29)	(0.07)	(-0.72)	(0.66)
$\text{Mkt} - \text{RF}_t$	0.9588***	0.9105***	0.8943***	0.8905***	0.9344***
	(58.27)	(57.10)	(59.02)	(59.60)	(53.62)
SMB_t	0.7504***	0.6753***	0.5366***	0.5572***	0.7497***
	(29.96)	(27.84)	(24.36)	(23.52)	(29.53)
HML_t	-0.0104	0.1618***	0.2083***	0.1447***	-0.0290
	(-0.37)	(6.41)	(8.67)	(6.18)	(-1.02)
MOM_t	-0.0954^{***}	-0.0993^{***}	-0.0936^{***}	-0.0907^{***}	-0.1261^{***}
	(-7.38)	(-8.26)	(-9.19)	(-8.39)	(-9.43)
RMW_t	-0.1876^{***}	-0.0226	-0.0440^{*}	-0.0386	-0.1867^{***}
	(-5.29)	(-0.77)	(-1.84)	(-1.44)	(-5.36)
CMA_t	0.0533	0.0545*	0.0981***	0.1069***	0.1155***
	(1.52)	(1.89)	(3.86)	(3.91)	(3.30)
Observations	101,394	112,462	123,171	121,207	103,878
# Stock	886	885	885	885	886
Stock FE	Yes	Yes	Yes	Yes	Yes
Within R^2	0.1449	0.1720	0.1926	0.1705	0.1522

The table reports estimates of stock-level exposure to innovations in semiconductor technology for portfolios sorted on their sensitivity to semiconductor producer prices. The sample consists of monthly stock returns from 1985 to 2024. I first split the sample at January 2005. Using the pre-2005 period, I estimate each stock's exposure to fluctuations in chip prices from the following time-series regression:

$$r_{i,t} = \alpha_i + \beta_i^{\text{ppi}}\,\Delta(\text{PPI}_t^{\text{semi}}) + \boldsymbol{\beta}_i^{\text{c}}\boldsymbol{C}_t + \varepsilon_{i,t},$$

where $r_{i,t}$ denotes the monthly excess return of stock i, $\Delta(\text{PPI}_t^{\text{semi}})$ is the monthly log growth of the producer price index for semiconductor and other electronic component manufacturing, and $\boldsymbol{C}_t$ includes the Fama–French five factors and the momentum factor. Based on the estimated pre-2005 values of β_i^{ppi}, stocks are assigned to five groups in January 2005. The post-2005 sample is then used to estimate, separately for each group, the exposure of stock returns to changes in the semiconductor innovation index:

$$r_{i,t} = \alpha_i + \beta^{\text{semi}}\,\Delta(\text{semi_idx}_t) + \boldsymbol{\beta}^{\text{c}}\boldsymbol{C}_t + \varepsilon_{i,t}.$$

The table reports the coefficient estimates from these post-2005 regressions. Stock fixed effects are included in all specifications. Standard errors are clustered at the stock level, and the corresponding t-statistics are reported in parentheses.

Table 10: Median β^{semi} from Pre-Ranking Regression, by Size, Market β, and Book-to-Market Quintiles

	Market Beta (β^{mkt})				
Size (market cap)	Low	2	3	4	High
Small	-0.04	-0.00	0.04	0.21	0.56
2	-0.12	-0.01	-0.02	0.14	0.26
3	-0.14	-0.06	-0.04	0.07	0.13
4	-0.11	-0.02	0.01	0.02	0.08
Big	-0.24	-0.03	0.01	0.14	0.30

	Book-to-Market				
Size (market cap)	Low	2	3	4	High
Small	-0.04	0.04	0.26	0.20	0.05
2	-0.11	-0.25	0.12	0.16	0.15
3	-0.11	0.02	0.03	0.12	0.19
4	-0.05	-0.12	0.03	0.19	0.07
Big	-0.17	-0.06	0.19	0.04	0.04

The table above reports the median value of pre-ranking factor loadings for quintiles constructed based on stock's market beta, market capitalization, and book-to-market. To estimate the factor loadings, for each stock-month observation, I run the time-series regression below over the rolling window of the past 12 months, and collect monthly β^{innov}. The table above reports the distribution of β^{inv} estimated across all months and stocks.

$$r_{i,t} = \alpha_{i,t} + \beta_{i,t}^{\text{semi}}\,\Delta(\text{semi_idx}_\text{t}) + \beta_{i,t}^{\text{mkt}}\,\text{MKT}_t + \varepsilon_{i,t}.$$

In this equation, $r_{i,t}$ is the monthly excess returns over the risk-free rate for stock i. The variable semi_idx_shock_t represents the unexpected shocks to my innovation index at month t. The value of the shock series at month t is the last residual collected from an AR(1) model applied to $semi_idx$ series up to the current month. The shock series start after at least 30 observations of semi_idx_t series that itself starts from January 1981.

Table 11: Second Stage Regressions Results

	25 portfolios (5x5)				100 portfolios (10x10)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
β^{semi}	**-0.2322*****	**-0.2985*****	**-0.1772*****	**-0.1762*****	**-0.0814****	**-0.0922*****	**-0.0781*****	**-0.0658*****
	(-2.69)	(-4.39)	(-3.27)	(-3.42)	(-2.23)	(-3.24)	(-3.25)	(-2.79)
Num. Periods	442	442	442	442	442	442	442	442
Avg. adj R^2	0.0442	0.3694	0.4849	0.4933	0.0099	0.2065	0.2590	0.2637
Controls	–	Market Factor	FF3 + Mom	FF5 + Mom	–	Market Factor	FF3 + Mom	FF5 + Mom

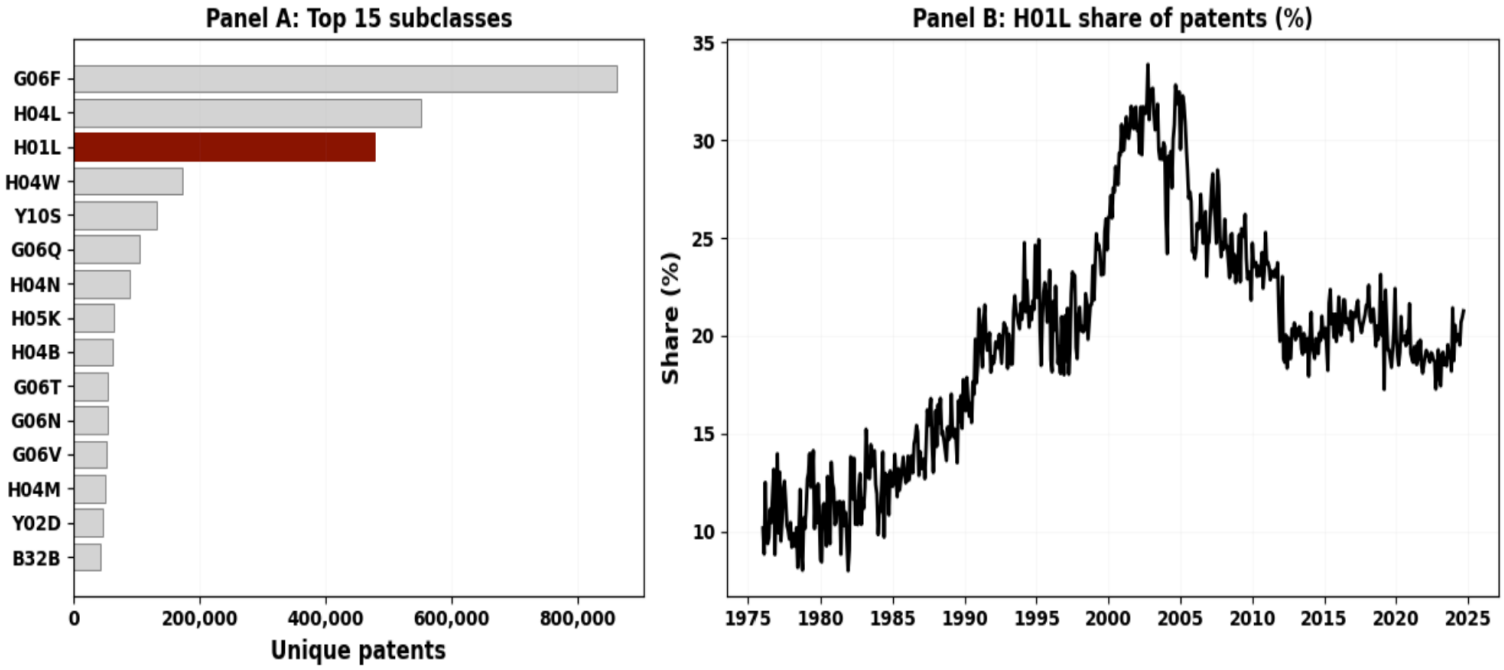
In the second stage regressions, I run the cross-sectional regression below for each month t separately. The table above reports the average of coefficients from this regression across all months, representing the price of risk associated with each factor.

$$r_{p,t} = \alpha_t + \boldsymbol{\pi_t^{\text{semi}}}.\boldsymbol{\beta_{p,t}^{\text{semi}}} + \pi_t^{\text{mkt}}.\beta_{p,t}^{\text{mkt}} + \pi_t^{\text{smb}}.\beta_{p,t}^{\text{smb}} + \pi_t^{\text{hml}}.\beta_{p,t}^{\text{hml}} + \\ + \pi_t^{\text{rmw}}.\beta_{p,t}^{\text{rmw}} + \pi_t^{\text{cma}}.\beta_{p,t}^{\text{cma}} + \pi_t^{\text{umd}}.\beta_{p,t}^{\text{umd}} + \varepsilon_{p,t}$$

where $p = 1, ..., 25$ for Columns (1) to (4) and $p = 1, ..., 100$ for Columns (5) to (8), representing the twenty five and hundred equal-weighted portfolios double-sorted based on pre-ranking factor loadings on excess market return and my innovation index. $r_{p,t}$ is the monthly excess returns over the risk-free rate for portfolio p in month t. $\beta_{p,t}$ factor loadings are estimated for each portfolio at month t on an expanding window of its past returns, using the first-stage regressions. Controls represent the factors included in both first and second-stage regressions. Reported in parentheses are Newey–West $t - statistics$ computed using average and standard errors of the time series for coefficients estimated from the cross-sectional regression above, across all months. Adjusted R^2 is the average adjusted R^2 from cross-sectional regressions across all months. In all specifications, stock-month observations with returns in the top and bottom 1% of the distribution are excluded before the construction of equal-weighted portfolios. The months used for the second-stage cross-sectional regressions span from January 1988 to October 2024.

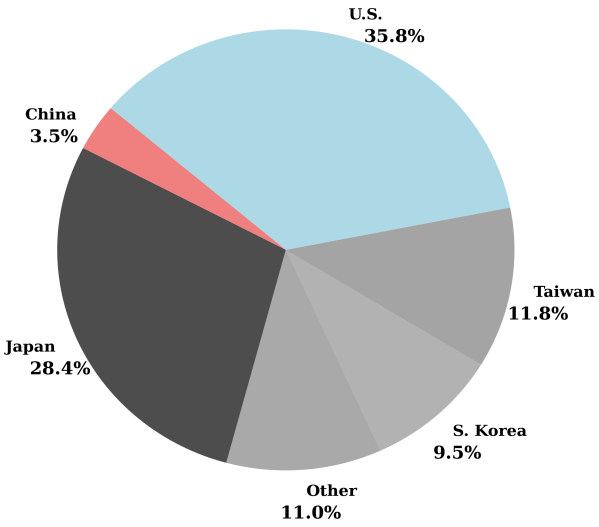
Figures

Figure 1: Share of Semiconductor Patents to Total Patents



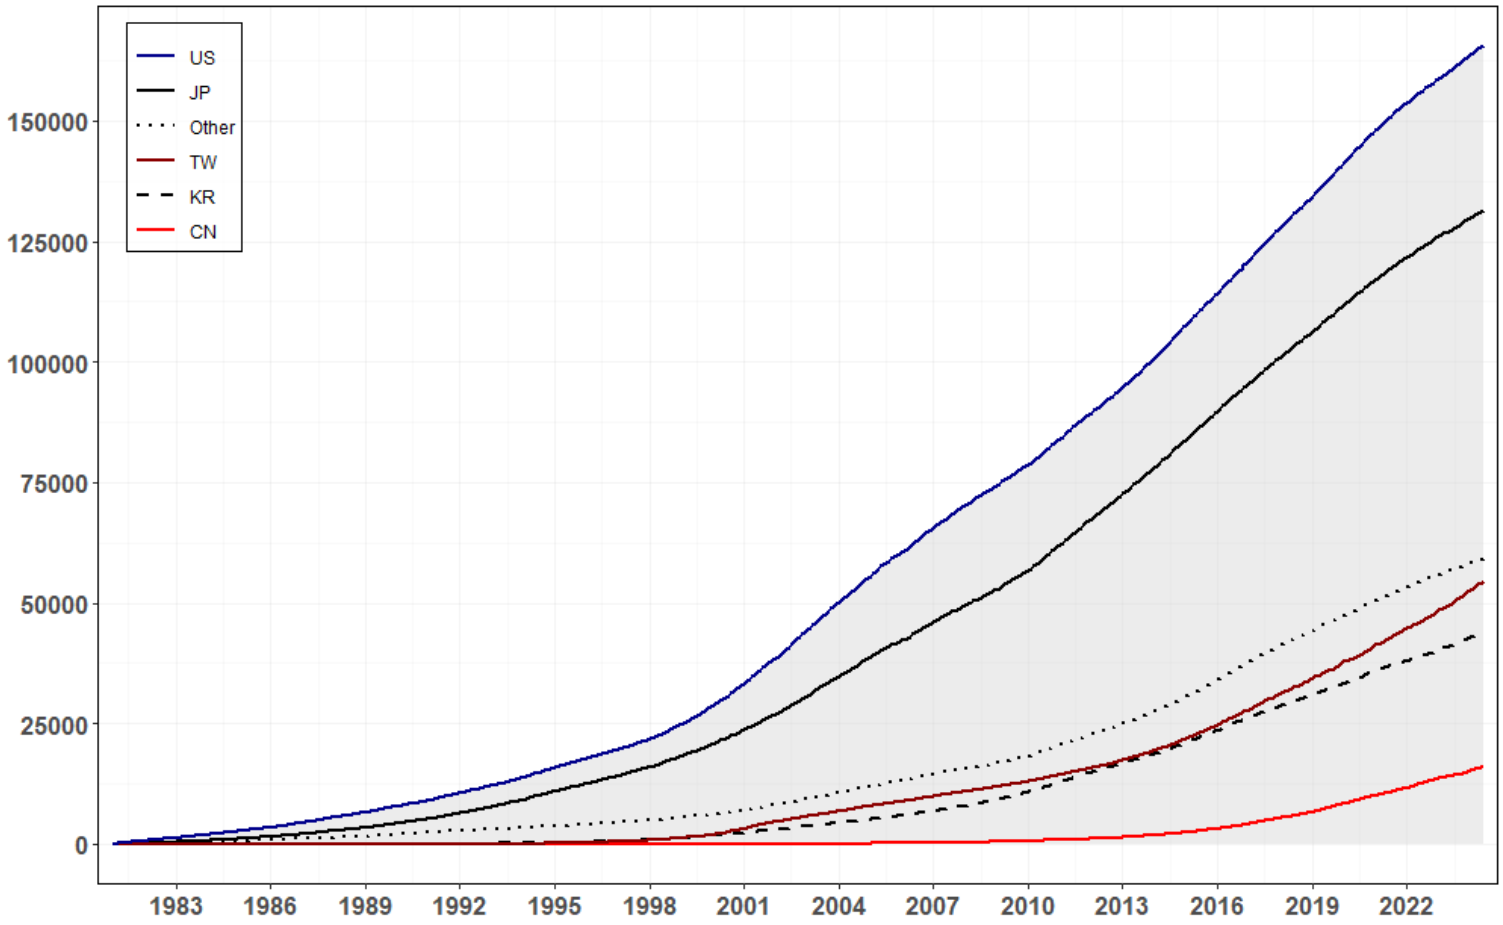
Panel A reports total unique patent counts by CPC subclass for the top 15 subclasses in the sample, with H01L highlighted. Panel B reflects unique patents per subclass (a patent may contribute to more than one bar). The sample consists of U.S. patents with CPC classifications between 1976 and 2025. Panel B plots the monthly share of patents assigned to CPC subclass H01L (Semiconductor Devices), computed as the number of unique patents tagged with H01L divided by the total number of unique patents with any CPC subclass in that month. I identify a total of 672 CPC subclasses in my sample of 8.1 million granted patents.

Figure 2: Top Countries with Granted Semiconductor Patents



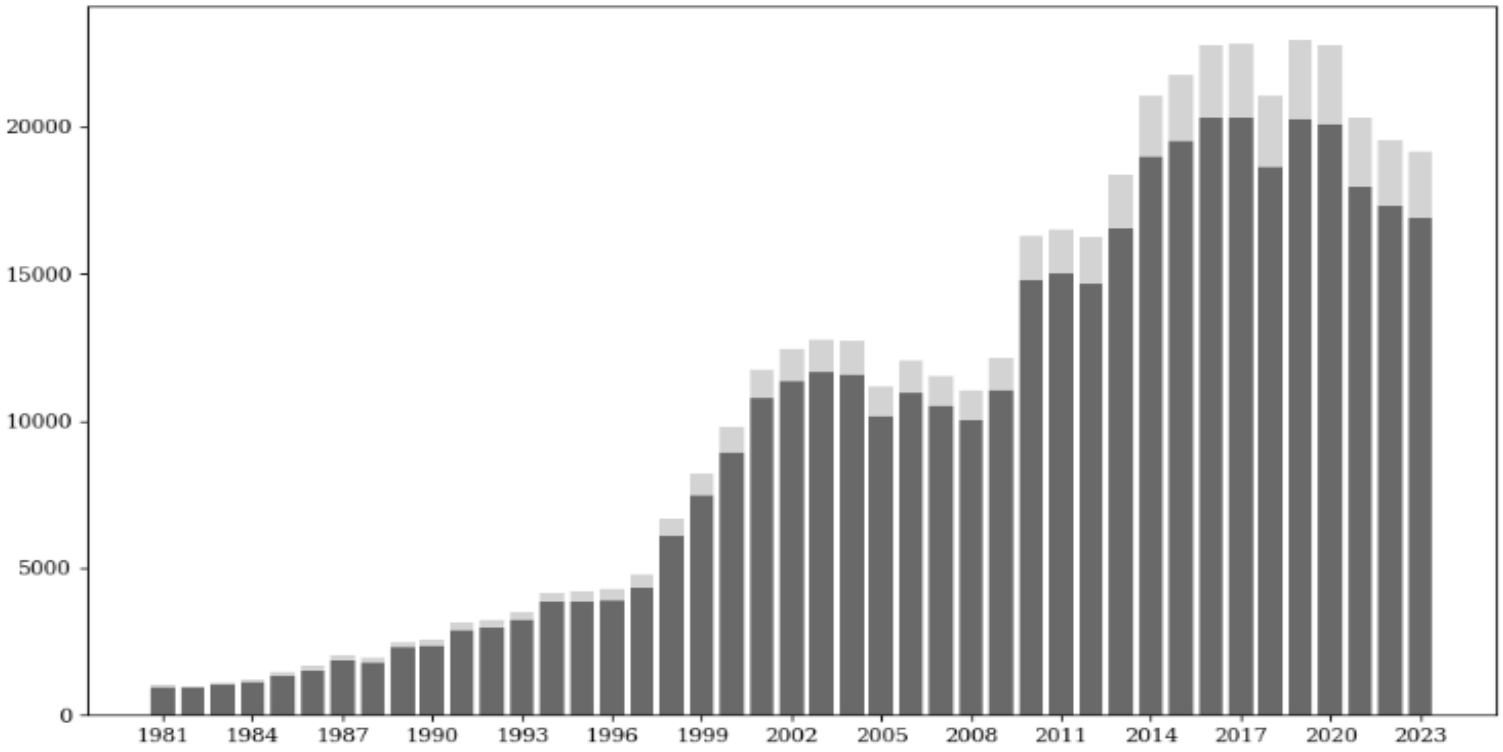
This pie chart illustrates the distribution of granted semiconductor patents by country for the period that starts from 1976 and ends in 2023. A patent is flagged for a country if that country code appears in the assignee list. Collaborative patents with multiple assignee countries contribute to each applicable country's series. Patents with none of {US, JP, TW, KR, CN} are counted under Other. Patents are categorized as semiconductor-related patents based on the 4-digit Cooperative Patent Classification (CPC) code.

Figure 3: Cumulative Number of Semiconductor Patents Granted by Country of Assignee



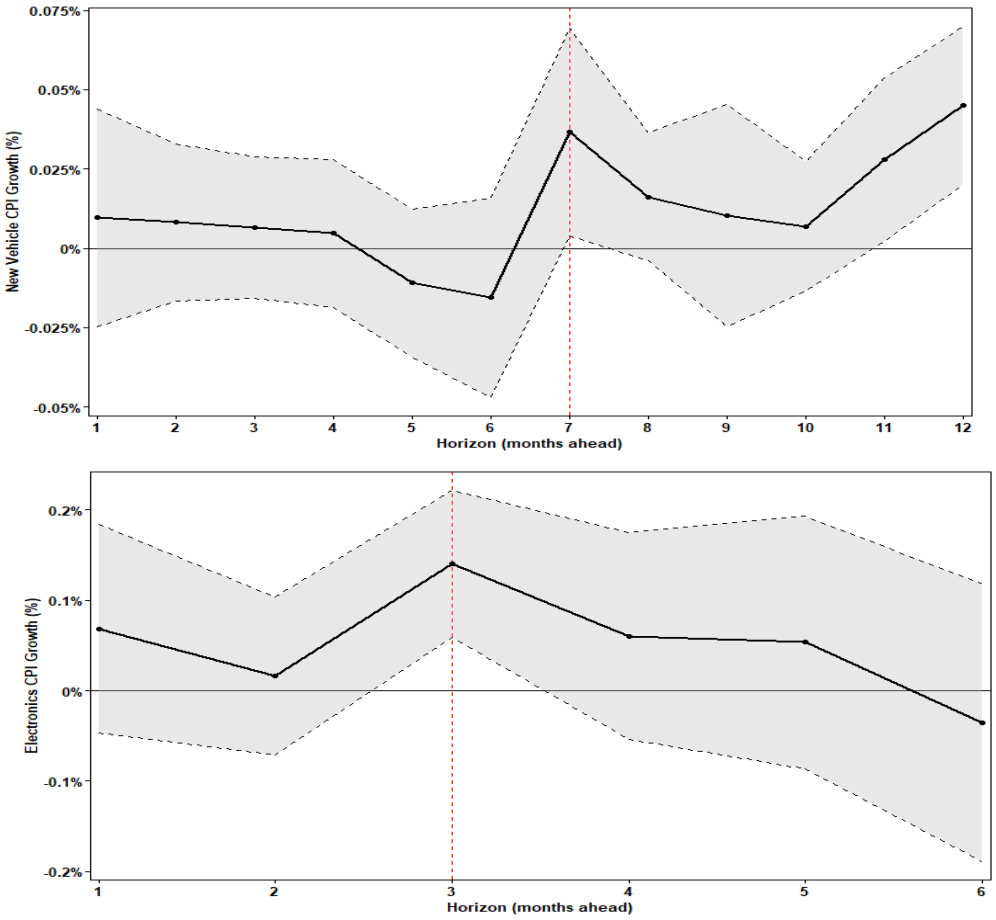
The figure plots the monthly cumulative number of patents granted by country of assignee from 1976–2023 for the top countries, including the U.S. (US), Japan (JP), Taiwan (TW), South Korea (KR), China (CN), and Other group. A patent is flagged for a country if that country code appears in the assignee list. Collaborative patents with multiple assignee countries contribute to each applicable country's series. Patents with none of {US, JP, TW, KR, CN} are counted under Other. Patents are categorized as semiconductor-related patents based on the 4-digit Cooperative Patent Classification (CPC) code.

Figure 4: Semiconductor Patent Count by Year



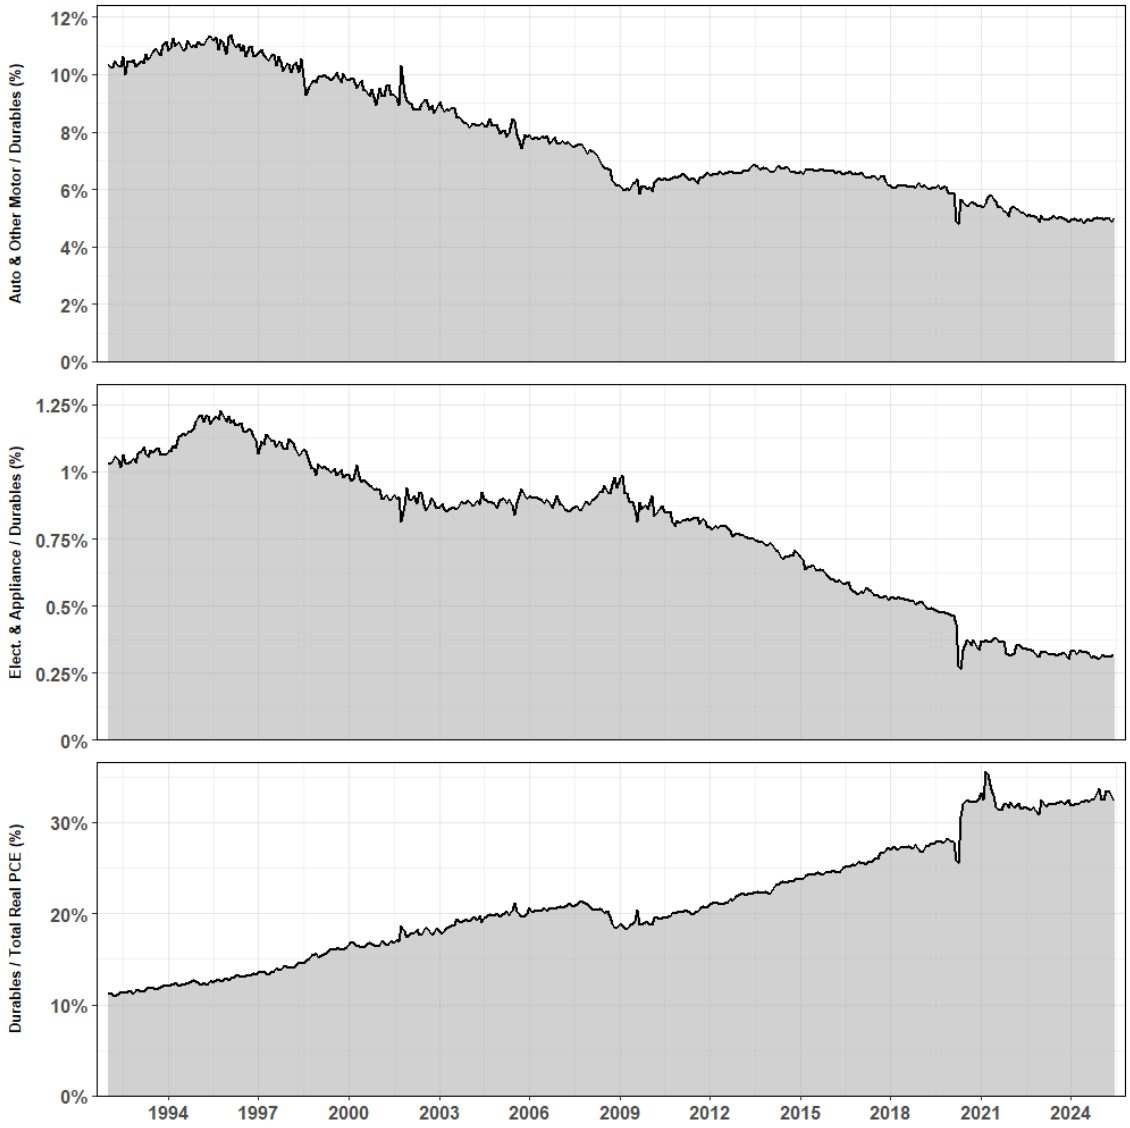
The figure reports annual counts of granted semiconductor patents from 1981 to 2023 (grant year on the horizontal axis). Each bar gives the total number of patents granted in that year, and the light segment at the top marks the subset of truly novel patents that fall in the top decile (90th percentile and above) by year. Patents are categorized as semiconductor-related patents based on the 4-digit Cooperative Patent Classification (CPC) code.

Figure 5: Durables Consumer Price Index Response to Producer Price Index for Semiconductors



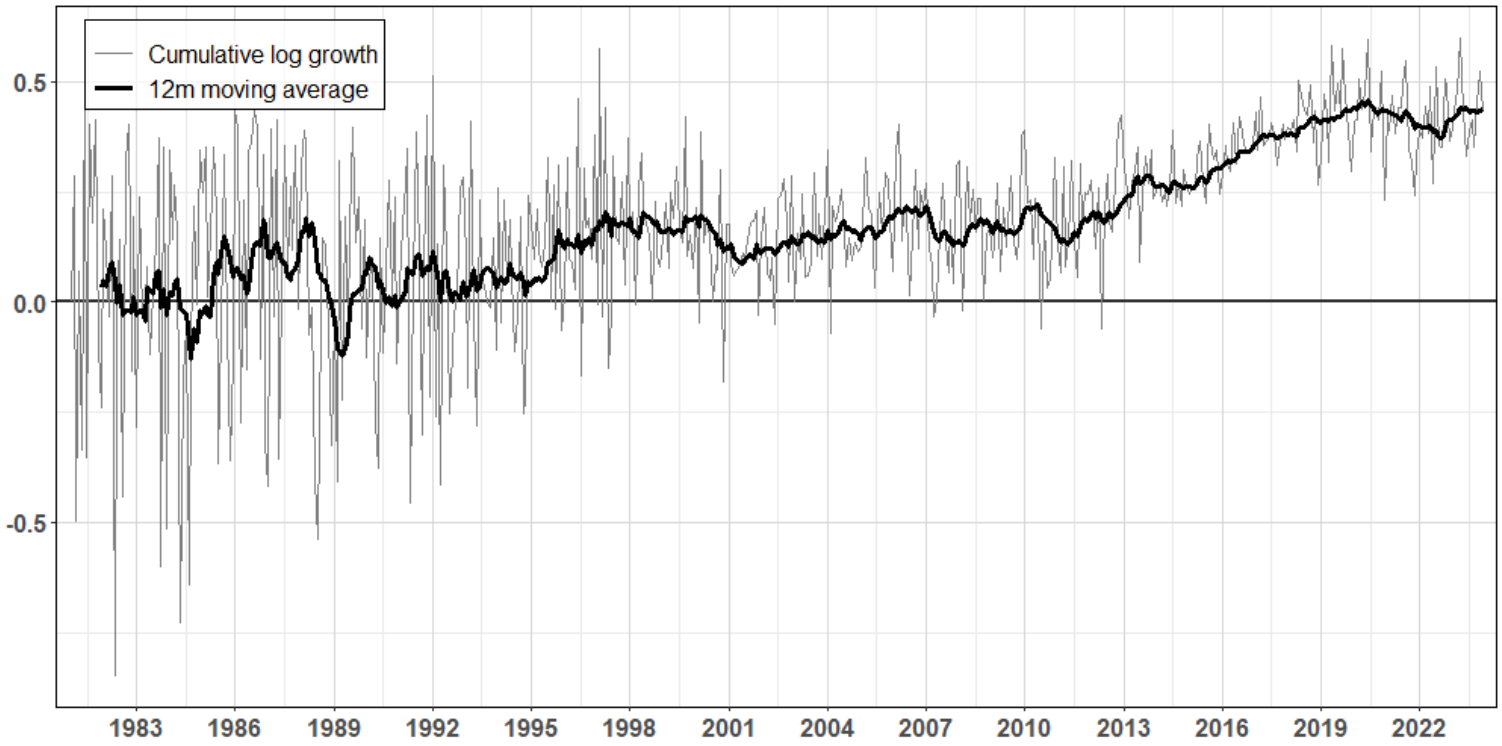
The figures report impulse response functions (IRFs) estimated via Local Projections (Jordà, 2005) at monthly horizons. The shock variable is the monthly log growth PPI^{semi}, standardised to one standard deviation using its standard deviation. Prior to estimation, PPI^{semi} is purged of predictable variation by regressing it on the full set of controls and retaining the residual, so that the identifying shock is orthogonal to the predictors of the dependent variables. The top figure reports the response of the log growth rate of the *Consumer Price Index for All Urban Consumers: New Vehicles in U.S. City Average* ($\Delta CPI_t^{vehicle}$), and the bottom figure reports the response for the *Consumer Price Index for All Urban Consumers: Computers, Peripherals, and Smart Home Assistants in U.S. City Average* ($\Delta CPI_t^{electronics}$). Both series are obtained from the Federal Reserve Bank of St. Louis's FRED database. The controls included in impulse response estimations for $\Delta CPI_t^{vehicle}$ are the same as in Table4 and for $\Delta CPI_t^{electronics}$ are the same as in Table4. All responses are expressed in percentage points. Shaded bands and dashed lines denote 95% confidence intervals constructed using Newey–West standard errors.

Figure 6: Retail Sales: Auto and Other Motor Vehicle and Retail Sales: Electronics and Appliance as a share of Personal Consumption Expenditure, Durable Goods



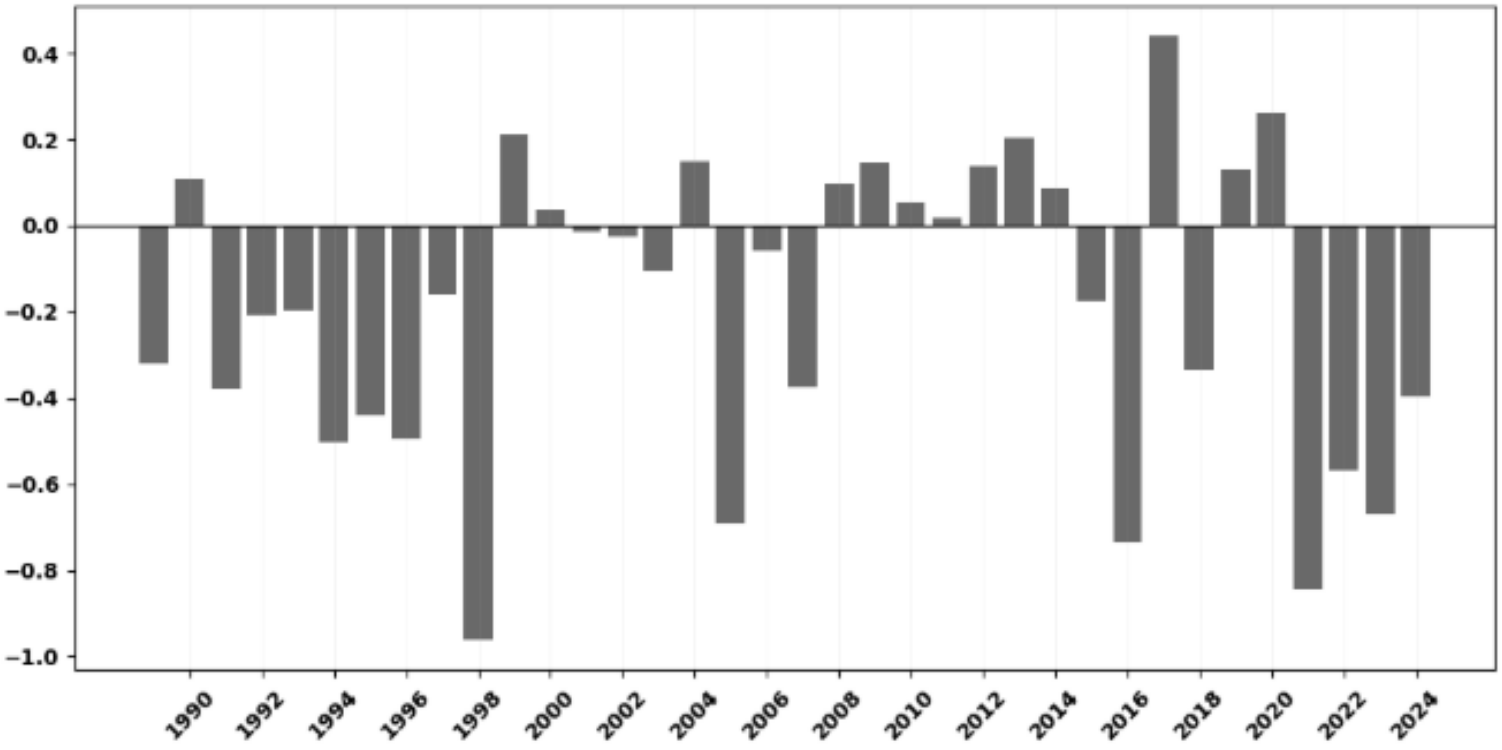
Top plot shows the ratio of *Retail Sales: Auto and Other Motor Vehicle Dealers* (deflated by the CPI for new vehicles) over *Personal Consumption Expenditure: Durable Goods*. The middle plot shows the ratio of *Retail Sales: Electronics and Appliance* (deflated by the Consumer Price Index (CPI) for durables) over *Personal Consumption Expenditure: Durable Goods*. Bottom plot shows the ratio of *Personal Consumption Expenditure: Durable Goods* over real total *Personal Consumption Expenditure*. All series are deflated and seasonally adjusted and are obtained from the Federal Reserve Bank of St. Louis's FRED database.

Figure 7: Cumulative Monthly Growth in Semiconductor Innovation Index



The figure shows the cumulative monthly log growth in the semiconductor innovation index ($\Delta(\text{semi_idx}_t)$) from 1981 to 2023. The semiconductor innovation index is defined as the share of patents that fall within the top 10th percentile of the novelty distribution—truly novel patents—over total patents granted each month. The plot also includes the 12-months moving average of the cumulative series to illustrate the persistent growth in semiconductor innovations over the years.

Figure 8: Price of Risk Associated with Semiconductor Innovations by Year



This figure presents the annual averages of the estimated price of risk (π_{innov}), as defined in Eq. 9. The sample period extends from January 1989 to October 2024.

References

Abernathy, William J. 1975. "A dynamic model of process and product innovation." *Omega* 3:639–656.

Abrams, David S, Ufuk Akcigit and Jillian Popadak. 2013. "Understanding the link between patent value and citations: Creative destruction or defensive disruption." *NBER Working Paper* 19647.

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom and William Kerr. 2018. "Innovation, reallocation, and growth." *American Economic Review* 108(11):3450–3491.

Adner, Ron and Daniel Levinthal. 2001. "Demand heterogeneity and technology evolution: implications for product and process innovation." *Management science* 47(5):611–628.

ASML. 2026. "EUV Lithography Systems.". Accessed: 2026-04-01.
URL: *https://www.asml.com/products/euv-lithography-systems*

Basu, Susanto and John Fernald. 2001. Why is productivity procyclical? Why do we care? In *New developments in productivity analysis*. University of Chicago Press pp. 225–302.

Basu, Susanto, John G Fernald and Miles S Kimball. 2006. "Are technology improvements contractionary?" *American Economic Review* 96(5):1418–1448.

Bloom, Nicholas, Raffaella Sadun and John Van Reenen. 2012. "Americans do IT better: US multinationals and the productivity miracle." *American Economic Review* 102(1):167–201.

Breeden, Douglas T. 1979. "An intertemporal asset pricing model with stochastic consumption and investment opportunities." *Journal of financial Economics* 7(3):265–296.

Breeden, Douglas T, Michael R Gibbons and Robert H Litzenberger. 1989. "Empirical tests of the consumption-oriented CAPM." *The Journal of Finance* 44(2):231–262.

Breitung, Christian and Sebastian Müller. 2025. "Global Business Networks." *Journal of Financial Economics* 166:104007.

Carhart, Mark M. 1997. "On persistence in mutual fund performance." *Journal of Finance* 52(1):57–82.

Cohen, Lauren, Christopher Malloy and Quoc Nguyen. 2020. "Lazy prices." *The Journal of Finance* 75(3):1371–1415.

Crespi, Gustavo, Chiara Criscuolo, Jonathan E Haskel and Matthew Slaughter. 2008. Productivity growth, knowledge flows, and spillovers. Technical report National Bureau of Economic Research.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*. pp. 4171–4186.

Fama, Eugene F and James D MacBeth. 1973. "Risk, return, and equilibrium: Empirical tests." *Journal of Political Economy* 81(3):607–636.

Fama, Eugene F and Kenneth R French. 2015. "A five-factor asset pricing model." *Journal of Financial Economics* 116(1):1–22.

Farre-Mensa, Joan, Deepak Hegde and Alexander Ljungqvist. 2020. "What is a patent worth? Evidence from the US patent "lottery"." *The Journal of Finance* 75(2):639–682.

Francis, Neville and Valerie A Ramey. 2005. "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited." *Journal of Monetary Economics* 52(8):1379–1399.

Gomes, Joao F, Leonid Kogan and Motohiro Yogo. 2009. "Durability of output and expected stock returns." *Journal of Political Economy* 117(5):941–986.

Graeber, Thomas, Christopher Roth and Florian Zimmermann. 2024. "Stories, statistics, and memory." *The Quarterly Journal of Economics* 139(4):2181–2225.

Greenwood, Jeremy, Zvi Hercowitz and Per Krusell. 1997. "Long-run implications of investment-specific technological change." *The American economic review* pp. 342–362.

Griliches, Zvi. 1961. Hedonic price indexes for automobiles: An econometric of quality change. In *The price statistics of the federal goverment*. NBER pp. 173–196.

Griliches, Zvi and Frank R Lichtenberg. 1984. R&D and productivity growth at the industry level: is there still a relationship? In *R&D, patents, and productivity*. University of Chicago Press pp. 465–502.

Haddad, Valentin, Paul Ho and Erik Loualiche. 2022. "Bubbles and the Value of Innovation." *Journal of Financial Economics* 145(1):69–84.

Hall, Bronwyn H and Rosemarie Ham Ziedonis. 2001. "The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979-1995." *rand Journal of Economics* pp. 101–128.

Hirshleifer, David, Po-Hsuan Hsu and Dongmei Li. 2013. "Innovative efficiency and stock returns." *Journal of financial economics* 107(3):632–654.

Hirshleifer, David, Po-Hsuan Hsu and Dongmei Li. 2018. "Innovative originality, profitability, and stock returns." *The Review of Financial Studies* 31(7):2553–2605.

Hoberg, Gerard and Gordon Phillips. 2010. "Product market synergies and competition in mergers and acquisitions: A text-based analysis." *The Review of Financial Studies* 23(10):3773–3811.

Hoberg, Gerard and Gordon Phillips. 2016. "Text-based network industries and endogenous product differentiation." *Journal of political economy* 124(5):1423–1465.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review* 95(1):161–182.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru and Matt Taddy. 2021. "Measuring technological innovation over the long run." *American Economic Review: Insights* 3(3):303–320.

Klepper, Steven. 1996. "Entry, exit, growth, and innovation over the product life cycle." *The American economic review* pp. 562–583.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru and Noah Stoffman. 2017. "Technological innovation, resource allocation, and growth." *The quarterly journal of economics* 132(2):665–712.

Kogan, Leonid, Dimitris Papanikolaou and Noah Stoffman. 2020. "Left behind: Creative destruction, inequality, and the stock market." *Journal of Political Economy* 128(3):855–906.

Kung, Howard and Lukas Schmid. 2015. "Innovation, growth, and asset prices." *The Journal of Finance* 70(3):1001–1037.

Lucas Jr, Robert E. 1978. "Asset prices in an exchange economy." *Econometrica: journal of the Econometric Society* pp. 1429–1445.

Mankiw, N Gregory and Matthew D Shapiro. 1984. Risk and return: Consumption versus market beta. Technical report National Bureau of Economic Research.

Mann, William. 2018. "Creditor rights and innovation: Evidence from patent collateral." *Journal of Financial Economics* 130(1):25–47.

Newey, Whitney K and Kenneth D West. 1987. "Hypothesis testing with efficient method of moments estimation." *International Economic Review* pp. 777–787.

Nicholas, Tom. 2008. "Does innovation cause stock market runups? Evidence from the great crash." *American Economic Review* 98(4):1370–1396.

Pakes, Ariel. 1985. "On patents, R & D, and the stock market rate of return." *Journal of political economy* 93(2):390–409.

Reimers, Nils and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. pp. 3982–3992.

Tham, Wing Wah, Salomé Baslandze, Elvira Sojli and Leo Liu. 2026. "Foundational processes and growth." *Journal of Financial Economics* 180:104276.

Yogo, Motohiro. 2006. "A consumption-based explanation of expected stock returns." *The Journal of Finance* 61(2):539–580.

Ziedonis, Rosemarie Ham. 2004. "Don't fence me in: Fragmented markets for technology and the patent acquisition strategies of firms." *Management science* 50(6):804–820.

Online Appendix
A1. Patent Abstracts

Table A1: Summary Statistics: Patent Abstract Length for Semiconductor Patents

Metric	N	Mean	SD	P10	P25	Median	P75	P90
Word count	472,724	114.8	41.78	62	85	115	143	158
Sentence count	472,724	3.53	2.19	1	2	3	5	7
Words per sentence	472,724	46.23	35.5	18.33	23.33	33	55.5	98

This table reports summary statistics for patent abstracts in the analysis sample. Word count is the number of word tokens in the abstract. Sentence count is the number of sentences identified in the abstract. Words per sentence is the ratio of word count to sentence count. Percentiles (P10, P25, P75, P90) are computed across abstracts.

A2. Semiconductor Innovation and Total Factor Productivity (TFP) Growth

To assess whether the semiconductor innovation index captures economically meaningful technological progress, we examine its predictive power for future productivity growth across industries. The rationale is that genuine innovation should translate into higher efficiency and output over time, as technological advances embodied in semiconductor components diffuse through production networks and enable improvements in downstream sectors. This approach aligns with extensive evidence showing that periods of heightened innovative activity are typically followed by sustained gains in total factor productivity, as innovation drives reallocation of resources toward more efficient technologies and firms (Kogan et al., 2017; Acemoglu et al., 2018). Following this logic, we test whether periods of heightened semiconductor innovation, as measured by our index, predict a rise in TFP. A positive association between innovation intensity and subsequent TFP growth would provide evidence that the index reflects innovations with real economic impact rather than strategic or financial patenting activity.

Given the pervasive role of semiconductors in modern production and technological systems, innovations in this sector are expected to influence productivity across a broad range of industries. However, the technology sector is expected to benefit most from semiconductor innovations because it both produces and intensively uses semiconductor-based technologies, which form the core of computing and information systems. As shown by **?**, productivity gains in IT-producing and IT-using industries, such as electronics, computing, and communication services, accounted for the bulk of the U.S. productivity acceleration since the mid-1990s. We therefore anticipate that periods of heightened semiconductor innovation will be followed by rises in TFP growth, with a larger effect concentrated in the technology sector.

To test this hypothesis, we examine the predictive power of the semiconductor innovation index for

subsequent productivity growth across industry groups. Specifically, we employ annual total TFP measures constructed at the 2- and 3-digit North American Industry Classification System (NAICS) levels, as reported by the U.S. Bureau of Labor Statistics. We use the panel regression below for our analysis:

$$\Delta \mathrm{TFP}_{i,t} = \alpha + \beta^{\mathrm{semi}} \Delta(\mathrm{semi_idx}_{[t-5:\,t-1]}) + \beta^{c}\, C_{i,t-1} + \delta_{ind} + \varepsilon_{i,t}. \tag{10}$$

Here, $\Delta \mathrm{TFP} i, t$ denotes the log change in TFP for industry i in year t, and $\Delta(\mathrm{semi_idx}_{[t-5:\,t-1]}$ represents the rolling cumulative log growth in the semiconductor innovation index over the preceding five years. We adopt this cumulative specification to capture the gradual and compounding nature of innovation's contribution to productivity. We also estimate specifications with $\Delta(\mathrm{semi_idx}_{[t-6:\,t-2]})$ and $\Delta(\mathrm{semi_idx}_{[t-7:\,t-3]})$ to account for the possible delayed effect of semiconductor innovation on TFP growth. to ensure that the estimated relationship is not merely driven by an increase in the overall volume of semiconductor patenting, we include $\Delta(\mathrm{total_s emi_p atents}_{[t-5:,t-1]})$, the rolling cumulative log growth in total semiconductor patent counts, as a control variable.

The vector of control variables, $C_{i,t-1}$, includes log growth in total factor productivity, output, capital input, and labor input in year $t-1$ for industry i. These controls are intended to account for cyclical fluctuations and other underlying dynamics that influence productivity. Lagged output growth is included to capture the procyclical nature of productivity, which tends to rise during booms and decline during recessions (Basu and Fernald, 2001). Lagged capital input growth controls for the potential impact of newly introduced, and potentially more efficient capital goods as an important source of productivity improvements (Greenwood, Hercowitz and Krusell, 1997). Lastly, lagged labor input growth is included to account for the potential effect of labor adjustments on productivity. The sample covers 81 sectors based on 2- and 3-digit, over the period between 1987 and 2023.

Table A2 presents panel regression estimates of the relationship between semiconductor innovation and subsequent growth in TFP across industries. Columns (1)–(6) show results for all sectors, while Columns (7)–(9) focus on industries in the technology sector.[13] Across specifications, the rolling 5-year cumulative growth in the semiconductor innovation index significantly predicts future TFP growth, with the effect materializing after a lag of roughly two to three years.

The results also reveal substantial sectoral heterogeneity. In Columns (1)–(6), the estimated coefficients in $\Delta(\text{semi}_\text{i}\text{dx}_{[t-7:t-3]})$ are positive and statistically significant, while for Columns (7)–(9), the estimated coefficients for $\Delta(\text{semi}_\text{i}\text{dx}_{[t-6:t-2]})$ and $\Delta(\text{semi}_\text{i}\text{dx}_{[t-7:t-3]})$ are both positive and statistically significant. In addition, the magnitude of the coefficient for $\Delta(\text{semi}_\text{i}\text{dx}_{[t-7:t-3]})$ estimated for industries in the technology sector is three times larger than its estimated value across industries from all sectors. These results indicate that the response of TFP is both earlier, emerging after roughly two years, and approximately three times larger in the technology sector relative to the economy-wide average. This pattern aligns with the sector's intensive use and production of semiconductor-based technologies, which allows it to absorb and convert innovation into productivity improvements more rapidly and effectively than other sectors. Importantly, the semiconductor innovation index's predictive power is not captured by the overall volume of semiconductor patenting. The coefficients on the 5-year rolling cumulative log growth in total semiconductor patent counts are small and statistically insignificant, suggesting that simple patent counts do not explain productivity variation once the novelty-based innovation index is accounted for.

Overall, the evidence indicates that semiconductor innovations are a significant predictor of productivity growth in the U.S. economy. The impact is particularly pronounced in the technology sector, where productivity responds earlier and with roughly three times the magnitude observed in the broader econ-

[13] The technology sector includes the following industries based on NAICS two-digit classification: Information (code 51), Professional, Scientific, and Technical Services (code 54).

omy. This pattern is consistent with the sector's greater dependence on semiconductor technologies and provides further validation that our semiconductor innovation index effectively captures meaningful variation in semiconductor innovation over time. Moreover, these results are not explained by mere increases in patenting activity, underscoring that the semiconductor innovation index captures technological innovation beyond mere patent count.

Table A2: Semiconductor Innovation and Total Factor Productivity

	Dependent variable: $\Delta TFP_{i,t}$:								
	Panel A (All Sectors)						*Panel B (Tech Sector)*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\sum_{i=1}^{5} \Delta(\text{semi_idx}_{[t-i]})$	−0.0347						−0.0103		
	(−0.88)						(−0.07)		
$\sum_{i=1}^{5} \Delta(\text{semi_idx}^{KPST}_{[t-i]})$		0.0722							
		(1.45)							
$\sum_{i=2}^{6} \Delta(\text{semi_idx}_{[t-i]})$			0.0177					0.1720	
			(0.41)					(1.88)	
$\sum_{i=2}^{6} \Delta(\text{semi_idx}^{kpst}_{[t-i]})$				0.1246					
				(1.75)					
$\sum_{i=3}^{7} \Delta(\text{semi_idx}_{[t-i]})$					**0.1207**				**0.4244**
					(2.70)				(3.67)
$\sum_{i=3}^{7} \Delta(\text{semi_idx}^{kpst}_{[t-i]})$						0.0960			
						(1.65)			
$\Delta Onput_{i,t-1}$	0.4818	0.4791	0.4814	0.4780	0.4854	0.4816	0.3884	0.3761	0.3635
	(7.56)	(7.40)	(7.61)	(7.32)	(7.60)	(7.33)	(2.70)	(2.82)	(2.87)
$\Delta Input^{capital}_{i,t-1}$	−0.1084	−0.1096	−0.1082	−0.1100	−0.1065	−0.1071	−0.1658	−0.1588	−0.1632
	(−1.99)	(−2.00)	(−2.00)	(−2.04)	(−1.97)	(−2.00)	(−3.32)	(−3.71)	(−4.23)
$\Delta Input^{labour}_{i,t-1}$	−0.3100	−0.3113	−0.3086	−0.3098	−0.3124	−0.3127	−0.4075	−0.4008	−0.3656
	(−5.07)	(−5.01)	(−5.05)	(−4.94)	(−5.07)	(−4.95)	(−2.07)	(−2.04)	(−1.93)
Observations	2,916	2,916	2,916	2,916	2,835	2,835	396	396	385
Adjusted R^2	0.314	0.313	0.314	0.313	0.317	0.315	0.113	0.112	0.132
Industry Fixed Effect	Y	Y	Y	Y	Y	Y	Y	Y	Y

The table reports estimates from the following annual industry-level panel regression:

$$\Delta \text{TFP}_{i,t} = \alpha + \beta^{\text{semi}} \sum_{i=1}^{5} \Delta(\text{semi_idx}_{[t-i]}) + \beta^{c}\, C_{i,t-1} + \delta_{ind} + \varepsilon_{i,t}.$$

Here, $\Delta \text{TFP}_{i,t}$ denotes the log growth in total factor productivity for industry i (2 and 3-digit NAICS) in year t, and $\sum_{i=1}^{5} \Delta(\text{semi_idx}_{[t-i]})$ is the distributed lags of annual log growth in semiconductor innovation index for lags 1 to 5. Similarly $\sum_{i=2}^{6} \Delta(\text{semi_idx}_{[t-i]})$ and $\sum_{i=3}^{7} \Delta(\text{semi_idx}_{[t-i]})$ are distributed lags including lags 2 to 6, and lags 3 to 7. For comparison, we also estimate specifications with $\sum_{i=1}^{5} \Delta(\text{semi_idx}^{KPST}_{[t-i]})$ representing the distributed lags of annual log growth in semiconductor innovation index, constructed based on $KPST$ patent novelty measure. For comparison, we also include $\Delta(\text{total_semi_patents}_{[t-5:\,t-1]})$ as the rolling cumulative log growth in the total semiconductor patents count in the past 5 years. $C_{i,t-1}$ is the vector of control variables, comprising the log growth in output, capital input, and labor input in year $t-1$ for industry i. Panel A includes industry observations across 81 industries from all sectors, while Panel B only includes observations from 11 industries in the technology sector, both covering the period between 1987 to 2023. The technology sector includes the following industries based on NAICS two-digit classification: Information (code 51), Professional, Scientific, and Technical Services (code 54). All regressions include industry fixed effects, and standard errors are clustered at the industry level. Year fixed effects are omitted as semi_idx varies only over time. t-statistics are reported in parentheses. Industry-level data are sourced from the US Bureau of Labor Statistics, tables for historical major sector total factor productivity (see Bureau of Labor Statistics).

Table A3: Semiconductor Innovation and Chip-Intensive Goods Retail Sales

	Dependent variable: $\Delta Retail\ Sales_t$:			
	Electronics & Appliance Stores		*Auto & Other Motor Dealers*	
	(1)	(2)	(1)	(2)
Constant	0.210	0.209	0.033	0.033
	(1.507)	(1.500)	(0.785)	(0.775)
$\Delta(\text{semi_idx}_{t-1})$	−0.023		0.008	
	(−0.607)		(0.840)	
$\Delta(\text{semi_idx}_{t-2})$		0.018		**−0.012****
		(0.785)		(−2.591)
$\Delta Retail\ Sales_{t-1}$ (*Electr. & Appliance*)	0.092**	0.092**		
	(2.948)	(3.102)		
$\Delta Retail\ Sales_{t-1}$ (*Auto & Motor*)			−0.126	−0.125
			(−1.515)	(−1.515)
$\Delta Disp.\ Income_{t-1}$	−0.045	−0.047	0.267	0.268
	(−0.435)	(−0.464)	(1.745)	(1.762)
Δ Comp. CPI_{t-1}	−0.783*	−0.756*		
	(−1.996)	(−2.118)		
Auto Loan Rate_{t-1}			0.0001	0.0002
			(0.070)	(0.080)
Δ Gas. CPI_{t-1}			−0.050	−0.052
			(−0.763)	(−0.792)
$log(WTI_{t-1})$	−0.025	−0.024	−0.011*	−0.011*
	(−1.629)	(−1.619)	(−2.014)	(−2.002)
$FF\ Rate_{t-1}$	−0.001	−0.001	0.0001	0.0001
	(−0.434)	(−0.422)	(0.074)	(0.062)
VIX_{t-1}	−0.003	−0.003	−0.001	−0.001
	(−1.357)	(−1.349)	(−1.594)	(−1.591)
EPU_{t-1}	0.0002	0.0002	0.0002***	0.0002***
	(1.602)	(1.613)	(3.604)	(3.598)
Cons. Sentiment_{t-1}	−0.001	−0.001	−0.00002	−0.00002
	(−1.557)	(−1.558)	(−0.084)	(−0.072)
Observations	226	226	382	382
Adjusted R^2	0.093	0.092	0.078	0.080

The table reports estimates from the following monthly time series regression:

$$\Delta Retail\ Sales_t = \alpha + \beta_k^{\text{semi}}\,\Delta(\text{semi_idx}_{t-k}) + \beta^c\,C_{t-1} + \varepsilon_t.$$

$\Delta Retail\ Sales_t$ denotes the monthly log growth of real Retail Sales series for the following categories: *Electronic and Appliance Stores* (deflated by the Consumer Price Index (CPI) for durables), *Auto and Other Motor Vehicle Dealers* (deflated by the CPI for new vehicles), all seasonally adjusted and obtained from the Federal Reserve Bank of St. Louis's FRED database. $\Delta(\text{semi_idx}_{t-k})$ is the seasonally adjusted log growth of the semiconductor innovation index in month $t-k$ ($k \in \{1,2\}$). The control vector C_{t-1} includes the lag of the dependent variables, the lagged log growth of *Real Disposable Personal Income*; the lagged log growth of *Consumer Price Index for All Urban Consumers: Computers, Peripherals, and Smart Home Assistants in U.S. City Average*; the lagged log growth of *Consumer Price Index for All Urban Consumers: Gasoline (All Types) in U.S. City Average*; the lagged *CBOE Volatility Index (VIX)*; the lagged *Economic Policy Uncertainty (EPU)* index; the lagged *Federal Funds Effective Rate*; the lagged *University of Michigan's Consumer Sentiment index*. All the controls are obtained from the Federal Reserve Bank of St. Louis's FRED database. Reported in parentheses are Newey–West t-statistics. The sample covers 382 monthly observations from 1992 through 2023.

Table A4: Semiconductor Innovation and Motor Vehicles Retail Inventory and Shipment

	Dependent variable:			
	$\Delta Retail\ Inv_t$		$\Delta Manuf.\ Ship_t$	
	(1)	(2)	(1)	(2)
$Constant$	0.039*	0.039*	0.127	0.126
	(2.163)	(2.153)	(1.244)	(1.228)
$\Delta(\text{semi_idx}_{t-1})$	−0.002		−0.013	
	(−0.530)		(−1.184)	
$\Delta(\text{semi_idx}_{t-2})$		−0.0002		−0.012
		(−0.072)		(−0.756)
$\Delta Retail\ Inv_{t-1}$	0.519***	0.519***		
	(8.494)	(8.493)		
ΔVehicle Assemblies$_{t-1}$	0.014***	0.014***		
	(6.492)	(6.521)		
ΔTotal Vehicle$_{t-1}$	0.035	0.035		
	(1.901)	(1.905)		
$\Delta Manuf.\ Ship_{t-1}$			−0.362***	−0.372***
			(−4.226)	(−4.520)
Δ Ind. Prod$_{t-1}$			−1.330	−1.349
			(−0.843)	(−0.854)
Δ Cap Util$_{t-1}$			1.669	1.693
			(1.024)	(1.040)
Δ Gas. CPI$_{t-1}$	0.005	0.005	0.034	0.036
	(0.225)	(0.238)	(0.814)	(0.870)
$\log(\text{WTI})_{t-1}$	−0.003	−0.003	−0.021	−0.021
	(−1.173)	(−1.169)	(−1.596)	(−1.590)
$FF\ \text{Rate}_{t-1}$	−0.0001	−0.00005	−0.0001	−0.00005
	(−0.144)	(−0.136)	(−0.045)	(−0.027)
VIX_{t-1}	−0.0003*	−0.0003*	−0.003	−0.003
	(−2.016)	(−2.023)	(−1.585)	(−1.596)
EPU_{t-1}	−0.0001*	−0.0001*	0.0003	0.0003
	(−2.187)	(−2.182)	(1.890)	(1.935)
Cons. Sentiment$_{t-1}$	−0.0002	−0.0002	−0.0004	−0.0004
	(−1.898)	(−1.888)	(−0.753)	(−0.732)
Observations	382	382	382	382
Adjusted R^2	0.403	0.402	0.108	0.108

This table reports estimates from monthly time–series regressions, where dependent variables are $\Delta Retail\ Inv_t$ and $\Delta Manuf.\ Ship_t$ as monthly log growth of the seasonally adjusted *Retail Inventories: Motor Vehicle and Parts Dealers* and *Manufacturers' Value of Shipments: Motor Vehicles and Parts* respectively, obtained from the Federal Reserve Bank of St. Louis's FRED database. $\Delta(\text{semi_idx}_{t-k})$ is the seasonally adjusted log growth of the semiconductor innovation index in month $t-k$ ($k \in \{1,2\}$). The control vector includes the lag of the dependent variables, the lagged log growth of *Industrial Production: Manufacturing: Durable Goods: Motor Vehicles and Parts* (in millions of units), the lagged log growth of *Capacity Utilization: Manufacturing: Durable Goods: Motor Vehicles and Parts* (in percentage), the lagged log

growth of *Real Disposable Personal Income*; the lagged log growth of *Consumer Price Index for All Urban Consumers: Computers, Peripherals, and Smart Home Assistants in U.S. City Average*; the lagged log growth of *Consumer Price Index for All Urban Consumers: Gasoline (All Types) in U.S. City Average*; the lagged *CBOE Volatility Index (VIX)*; the lagged *Economic Policy Uncertainty (EPU)* index; the lagged *Federal Funds Effective Rate*; the lagged *University of Michigan's Consumer Sentiment index*. All the controls are obtained from the Federal Reserve Bank of St. Louis's FRED database. Reported in parentheses are Newey–West t-statistics. The sample covers 382 monthly observations from 1992 through 2023.

A4. Estimating Price of Risk

For each variable in regression Eq.9, let $\hat{\lambda}_t$ denote the estimated price of risk from the monthly cross-sectional regression at time t. The code computes the time-series average:

$$\bar{\lambda} = \frac{1}{T}\sum_{t=1}^{T}\hat{\lambda}_t$$

The standard error of $\bar{\lambda}$ is estimated using the Newey and West (1987) heteroskedasticity and autocorrelation consistent estimator with $L = 12$ lags:

$$\mathrm{SE}_{\mathrm{NW}}(\bar{\lambda}) = \sqrt{\frac{1}{T}\left(\gamma_0 + 2\sum_{\ell=1}^{L}\left(1 - \frac{\ell}{L+1}\right)\gamma_\ell\right)}$$

where γ_ℓ is the sample autocovariance of $\hat{\lambda}_t$ at lag ℓ and are estimated as follows:

$$\gamma_0 = \frac{1}{T}\sum_{t=1}^{T}\left(\lambda_t - \bar{\lambda}\right)^2$$

$$\gamma_\ell = \frac{1}{T}\sum_{t=\ell+1}^{T}\left(\lambda_t - \bar{\lambda}\right)\left(\lambda_{t-\ell} - \bar{\lambda}\right), \quad \ell = 1, 2, \ldots, L$$

The corresponding t-statistic is then given by:

$$t = \frac{\bar{\lambda}}{\mathrm{SE}_{\mathrm{NW}}(\bar{\lambda})}$$

The reported values in Table 11 are $\bar{\lambda}$ and t statistics calculated as above.