RIETI Discussion Paper Series 26-E-045

Predicting Shock Propagation and Uncovering Heterogeneity with Graph Neural Networks

Yoshiyuki ARATA
RIETI

Research Institute of Economy, Trade & Industry, IAA
The Research Institute of Economy, Trade and Industry
https://www.rieti.go.jp/en/

RIETI Discussion Paper Series 26-E-045
May 2026

Predicting Shock Propagation and Uncovering Heterogeneity with Graph Neural Networks*

Yoshiyuki Arata
Research Institute of Economy, Trade and Industry

Abstract

Recent research has made substantial progress in studying shock propagation through inter-firm transaction networks, and empirical studies have directly documented firm-level shock propagation. Despite these advances in both theory and empirics, no method has yet been established to accurately predict the effects of large-scale future shocks, such as natural disasters, financial crises, or pandemics. A central challenge is the heterogeneity inherent in firms and transaction relationships, which makes it difficult to identify which firms are important for shock propagation and which links amplify it. To address this issue, this study uses firm-level data and a graph neural network (GNN) to predict firm growth rates with a model that explicitly incorporates network structure. In particular, by analyzing the trained GNN model, we quantitatively identify the firms and transaction links that are important for shock propagation. Using the global financial crisis, specifically the sharp decline in exports, as a case study, we show that incorporating network structure significantly improves predictive performance and enables us to identify specific firms and links that are important for propagation.

Keywords: Graph Neural Networks; Propagation of Demand Shocks
JEL classification: D22

*This study is conducted as a part of the Project "Study Group on Corporate Finance and Firm Dynamics " undertaken at the Research Institute of Economy, Trade and Industry (RIETI). Arata gratefully acknowledges financial support from the Japan Society for the Promotion of Science (Grant Number: 26K04904)

1 Introduction

The idea that the propagation of idiosyncratic shocks through input-output networks is important for the economy as a whole has long been recognized in macroeconomics. Classical studies by Long and Plosser (1983), Horvath (1998), Horvath (2000) show that inter-industry linkages can serve as channels through which shocks propagate and thereby affect aggregate fluctuations. This perspective has been further developed in the recent production network literature by explicitly incorporating firm-level input-output networks. In particular, studies since Acemoglu et al. (2012) have theoretically shown that the network structure itself shapes the aggregate effects of shocks, and that shocks to firms located at the center of the network can have large effects on the overall economy. Furthermore, advances in empirical research using firm-level data have made it possible to directly observe such shock propagation on observed firm-to-firm networks. In addition, with growing interest in economic security and supply chain resilience in recent years, understanding shock propagation has become important not only academically but also from a policy perspective.

However, despite this progress in research, methods for accurately predicting how future large-scale shocks will propagate are still not sufficiently established. A fundamental reason is the heterogeneity of firms and trading relationships. Shocks do not propagate uniformly across trading partners; in some cases they spread strongly, while in others they do not. For example, in a demand shock such as the Financial Crisis, the extent of the impact differs greatly depending on whether firms can find alternative customers. Likewise, in a supply shock such as the Great East Japan Earthquake, the substitutability of supplied goods determines the scope of propagation. As Barrot and Sauvagnat (2016) also show, differences in input specificity and substitutability play a crucial role in shock propagation. However, it is not easy to directly estimate from firm-level data parameters such as the elasticity of substitution or the possibility of finding new buyers and suppliers. As a result, highly accurate simulations and predictions that take into account which firms and which trading relationships play important roles in shock propagation have not yet been achieved.

To address this problem, this study adopts a prediction approach using Graph Neural Networks (GNNs). Because GNNs can explicitly incorporate not only firm attributes but also the structure of the interfirm transaction network itself into the model, they allow predictions that reflect the hierarchical connectedness of trading relationships. In particular, the process of updating node representations using information from neighboring nodes corresponds naturally to the mechanism through which shocks propagate over transaction networks. Within this framework, this study employs a Graph Attention Network (GAT). GAT learns, in the form of attention weights, which links should be emphasized and to what extent when aggregating information from adjacent nodes. Thus, it has the advantage not only of improving predictive accuracy but also of identifying the firms and transaction links that are important for shock propagation through analysis of the trained model. GAT therefore enables both prediction and interpretation that reflect heterogeneity, while alleviating to some extent the black-box nature of deep learning. As a concrete application, this study examines how the sharp decline in exports during the Financial Crisis propagated to domestic non-exporting firms in Japan.

Recent production network literature has theoretically shown that firm-level shocks can generate aggregate fluctuations through production networks. Studies since Acemoglu et al. (2012) have clarified that the structure of production networks itself is important for the aggregate effects of shocks (see Carvalho (2014); Carvalho and Tahbaz-Salehi (2019) for surveys). Earlier work focused primarily on modeling supply shocks, but studies such as Herskovic et al. (2020) have also examined the propagation of demand shocks. In addition, while Acemoglu et al. (2012) show that the first-order aggregate effect of shocks is proportional to firm size (i.e., firm size is the sufficient statistics), subsequent studies such as Baqaee and Farhi (2019) have shown that aggregate effects can become nonlinear depending on substitutability and complementarity among goods. In other words, when the goods produced by affected firms are difficult to substitute for, the impact of the shock can be substantially amplified. Moreover, Huneeus (2020) suggests that, because frictions exist in link formation and adjustment, there are limits to understanding shock propagation using only simple static models.

Empirical studies using firm-level data also strongly support the importance of shock propagation. Boehm et al. (2019) and Carvalho et al. (2021) use exogenous shocks such as natural disasters and supply disruptions to show that their effects propagate through firm-to-firm networks. In particular, Barrot and Sauvagnat (2016) show that supplier shocks propagate to customer firms through interfirm transaction relationships and suggest that propagation effects differ according to input specificity. More recently, studies such as Peter et al. (2023) and Fujiy et al. (2024) have also made progress in quantitatively measuring the substitutability of intermediate inputs. What these studies consistently show is that shocks do not propagate uniformly, and that the strength of propagation differs greatly across firms and links. Therefore, the key issue is not simply to show whether shocks propagate, but to clarify through which links and firms they propagate more strongly. However, for both supply and demand shocks, it is not easy to structurally estimate and model all factors such as substitution elasticities, search costs, and network position. This is a major reason why highly accurate shock simulations that take heterogeneity in firms and trading relationships into account have still not been achieved.

The key feature of this study is that it addresses the problems of network structure and heterogeneity by constructing a prediction model based on Graph Neural Networks and using it to predict the future growth rate of each firm. In conventional regression analysis and standard machine learning approaches, network structure is often compressed into averages of neighboring firms or other aggregated features, making it easy to lose information about which trading partners matter and which higher-order paths are important. By contrast, GNNs can explicitly incorporate network structure and generate predictions that reflect the hierarchical connectedness of trading relationships. In particular, the Graph Attention Network used in this study aggregates information while assigning attention weights to each link, thereby learning in a data-driven manner which firms and links are important for shock propagation. Thus, it not only improves predictive accuracy but also makes it possible, through analysis of attention, to interpret which adjacent links are more important and which firms' shocks have larger effects.[1]

The focus of this study is how the negative demand shock to exporting firms, in the form of the sharp decline in exports associated with the Financial Crisis, propagated to domestic non-exporting firms in Japan. Specifically, we treat Japanese exporting firms as having experienced an exogenous demand shock and predict the sales growth rates of domestic non-exporting firms. As node features, we use firm attributes such as firm age,

[1] In recent years, studies have also begun to apply Graph Neural Networks to the prediction of supply chain disruption. For example, Yang et al. (2024) uses GNNs to predict firm-level sales under post-disaster supply chain disruption, and is a closely related study that connects shocks and firm performance prediction through GNNs. Wasi et al. (2024) also provides applications of GNNs in supply chain analytics and discusses their usefulness relative to conventional classical machine learning.

industry, location, financial information, and text information such as business descriptions. As edge features, we incorporate the distinction between whether a transaction relationship is with a supplier or a customer, as well as the presence of a capital relationship. We train the GAT model on a subset of domestic non-exporting firms and evaluate its out-of-sample predictive accuracy on the remaining firms in order to assess prediction performance for shock propagation. In addition, by analyzing the attention of the trained model, we identify the firms and trading relationships that are important for shock propagation.

Empirically, the trained model achieves an R^2 of 0.286 on the test data, representing a substantial improvement in predictive accuracy relative to the case without network information. This indicates that, even in predicting firm growth rates that are strongly influenced by idiosyncratic factors, information on trading partners and network structure is useful. Furthermore, when we examine the important links and trading partners identified by the model for a large domestic non-exporting firm, most of the trading partners judged to be important are customers rather than suppliers. This is consistent with the fact that the Financial Crisis was primarily a demand shock in Japan and suggests that the results of the attention analysis are economically interpretable.

The remainder of this paper is organized as follows. Section 2 presents the formulation of the model using Graph Attention Networks. Section 3 describes the data, the definition of the sample, the construction of the variables, and the empirical design. Section 4 reports the prediction results for the Financial Crisis and evaluates the model's out-of-sample performance. Section 5 concludes.

2 Model

Here, we explain Graph Neural Networks (GNNs), in particular Graph Attention Networks (GATs) proposed by Veličković et al. (2018). A key feature of GNNs is that they predict the future growth rate of each firm using not only that firm's own attributes but also information from firms directly or indirectly connected through the transaction network. In particular, node representations can be updated while reflecting hierarchical connections that include second- and third-order trading partners, and this updating process itself can be interpreted as an approximation of shock propagation. In GATs especially, the degree to which each link should be emphasized can be learned as an attention weight, which makes it possible to identify important propagation paths together with prediction.

We now provide the formulation of GAT. In this analysis, the inter-firm transaction network is represented as

$$G = (V, E), \quad |V| = N$$

where each node $i \in V$ is a firm and each edge $(i, j) \in E$ is a transaction relationship. The adjacency matrix indicates which firms are directly connected, and thus determines how far each node can aggregate information, that is, the reachable range of information propagation. In particular, if the graph is defined as directed, the model can also represent the direction in which shocks flow.[2]

Each firm is assigned a node feature vector $x_i \in \mathbb{R}^d$, whose components include firm-specific attributes such as sales, industry classification, past growth rates, location, and firm age. This x_i is used as the initial input to the GAT, namely $h_i^{(0)} = x_i$. In addition, each transaction relationship is assigned an edge feature vector $e_{ij} \in \mathbb{R}^k$, so that heterogeneity in transaction relationships can be represented through information such as whether the counterparty is a supplier or a customer, and whether a capital relationship exists. The prediction target is the next-period growth rate y_i of each firm, and in implementation the model is treated as node-level regression, where one predicted value is output for each firm from the node representation obtained at the final layer.

In GAT, each node is considered to have an internal representation that describes its state, and the representation of node i aggregates not only the features of node i itself but also information from other nodes connected by links. A key feature of the update from layer ℓ to layer $\ell + 1$ is given as follows.

$$h_i^{(\ell+1)} = \sigma\left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(\ell)} W^{(\ell)} h_j^{(\ell)}\right)$$

Here, $h_i^{(\ell)}$ is the representation of node i at layer ℓ, and its initial value is $h_i^{(0)} = x_i$. $\mathcal{N}(i)$ is the neighborhood of node i, and by including a self-loop, the firm's own information is also retained in neighborhood aggregation. This update first applies a shared linear

[2]While the original GAT of Veličković et al. (2018) does not take edge features into account, the extended version proposed by Brody et al. (2022), which is employed in this study, incorporates edge features.

transformation $W^{(\ell)}$ to each node feature, then aggregates neighborhood information not by a simple average but by a weighted average using attention $\alpha_{ij}^{(\ell)}$, and finally passes it through a nonlinear transformation σ to represent higher-order effects. By stacking these layer-wise updates, the model aggregates information from more distant nodes in the network and constructs the representation of node i.

The attention weight $\alpha_{ij}^{(\ell)}$ is computed for each pair of node i and its neighboring node j as follows.

$$\alpha_{ij}^{(\ell)} := \frac{\exp(s_{ij}^{(\ell)})}{\sum_{k\in\mathcal{N}(i)} \exp(s_{ik}^{(\ell)})}, \quad s_{ij}^{(\ell)} := \mathrm{LeakyReLU}\left(a^\top [W_h h_i^{(\ell)} \| W_h h_j^{(\ell)} \| W_e e_{ij}]\right)$$

It is normalized by softmax, that is, $\sum_{j\in\mathcal{N}(i)} \alpha_{ij}^{(\ell)} = 1$. $\alpha_{ij}^{(\ell)}$ is a weight that indicates how much node j is emphasized when updating node i. Therefore, a larger $\alpha_{ij}^{(\ell)}$ means that the information of firm j contributes more strongly to the prediction for firm i, allowing heterogeneity to be handled more flexibly than in models that average trading partners uniformly. In addition, using multi-head attention, which computes multiple attention heads in parallel, increases expressive power.

The representation $h_i^{(L)}$ obtained at the final layer for each firm is fed into a linear output layer to obtain the predicted sales growth rate $\hat{y}_i$ for firm i:

$$\hat{y}_i = W_{\text{out}} h_i^{(L)} + b_{\text{out}},$$

Using the predicted values $\hat{y}_i$ for the training nodes, the model is trained to minimize the mean squared error (MSE) loss:

$$\mathcal{L} = \frac{1}{|V_{\text{train}}|} \sum_{i\in V_{\text{train}}} (y_i - \widehat{y}_i)^2.$$

In this analysis, we assume that a large negative growth rate is assigned to a node that serves as the origin of a shock (i.e., an exporting firm), and that this information is transmitted to neighboring nodes through message passing in GAT. An update of one layer corresponds to the influence received from directly connected trading partners, while updates over two or more layers correspond to propagation through more distant, indirect trading partners. Therefore, stacking layers makes it possible to represent the spread of shock propagation. In addition, links with larger attention weights $\alpha_{ij}^{(\ell)}$ can be interpreted as more important paths in the propagation of shocks. Accordingly, the process of updating node representations can be understood as the process by which firm-specific attributes and influences transmitted from neighboring firms are combined and reflected in next-period growth rates. Furthermore, by comparing prediction accuracy, we can examine how effective the network information itself and the introduction of attention are for prediction.

3 Data

3.1 Data

This study uses firm-level credit information data from Tokyo Shoko Research (TSR). TSR is a private credit research company that conducts firm surveys at the request of clients such as financial institutions and trading companies, and stores the results in a database. The data are constructed by integrating interviews conducted by survey staff, publicly available information, and materials submitted by firms.

A key feature of this data is that it covers not only listed firms but also a wide range of unlisted firms and small and medium-sized enterprises. The database contains more than one million firms and therefore provides useful coverage for analyzing transaction networks across Japanese firms. The data also have a panel structure, making it possible to observe firm attributes and financial information such as establishment year, location, industry, sales, profits, and the number of employees over time. In addition, the TSR data include information on transaction relationships between firms. Specifically, major suppliers and customers are recorded for each firm, up to a maximum of 24 firms, and this information can be used to construct an interfirm transaction network. Using this interfirm relationship data, this study analyzes shock propagation through supply chains at the firm level.

3.2 Model Settings and Samples

Here, we describe the empirical design used in our analysis. In this analysis, we treat the sharp decline in exports associated with the Financial Crisis as an exogenous negative demand shock to exporting firms. This shock is observed as a decline in the sales of exporting firms, and we assume that it propagates to other firms through interfirm transaction relationships. Because this is a demand shock, propagation through the customer-side links should be particularly important from a theoretical perspective. The prediction target in this study is the sales growth rate of domestic non-exporting firms from 2008 to 2010, and we examine whether shocks to exporting firms propagated through transaction relationships.

The network used in the analysis is restricted to firms with sales of at least one billion yen as of 2008 and to transaction relationships among those firms. Under this criterion, the number of firms in the sample is 74,081, and the sales threshold is introduced in order to secure a certain level of economic importance and data quality. Including small firms would increase missing values and noise and could make it more difficult to identify shock propagation, so the analysis is limited to firms above a certain scale. The network analyzed in the analysis is constructed from these firms and their mutual transaction relationships, and the transaction links are treated as fixed, with no subsequent link changes considered. Thus, the analysis is designed to capture short-run shock propagation on the existing network.

Among these firms, the prediction target nodes used for training and evaluation are about 22,294 non-exporting firms with sales of at least three billion yen, which are split into training and test sets. This threshold is also intended to ensure a certain scale and measurement accuracy on the prediction side, and the parameters are learned so that the growth rates of these firms can be predicted as accurately as possible. Exporting firms (7,885 firms) and firms with sales below three billion yen are not removed from the network; instead, they remain as nodes that serve only for information propagation and form part of the message-passing paths.[3] The network itself is shared between training and test, and the only difference is the set of nodes on which prediction errors are evaluated. This design makes it possible to evaluate out-of-sample performance for the target firms while preserving the structure of the full network.

As inputs to the GAT model, the node features include past sales growth, attribute and financial information (e.g., firm age, industry, location, sales, the number of employees, the number of business establishments, and the number of factories), as well as text information such as business descriptions. In addition, for exporting firms, the growth rate at the time of the shock is also included as a node feature that reflects the intensity of the shock. Since our prediction targets are non-exporting firms, it should be noted that the shock-period growth rates of exporting firms are used only when aggregating information from trading

[3] At present, due to data availability, exporting firms are defined not by their export status as of 2008 but by the presence of an export flag during 2012-2014.

partners. The edge features include whether a transaction relationship is with a supplier or a customer, and whether a capital relationship exists. The outcome variable is the two-year log growth rate of firm sales from 2008 to 2010. To reduce the influence of extreme values, an outlier treatment is applied in which log growth rates below -0.8 are replaced with -0.8 and values above 0.4 are replaced with 0.4.

In the modeling stage, a GAT model is trained using a subset of non-exporting firms as training data, and its out-of-sample predictive accuracy is evaluated on the remaining firms as test data. We also compare the model with cases that do not use network information in order to examine the importance of transaction relationships for prediction. Furthermore, we identify the firms and transaction relationships that are important for shock propagation from the trained model, and clarify which links and firms are important as propagation paths.

3.3 Descriptive Statistics

First, in order to confirm the magnitude of the shock to exporting firms during the Financial Crisis period, we examine the distribution of the sales growth rates of exporting firms over the two years from 2008 to 2010. This growth rate is a variable that directly captures the decline after the Lehman shock. In addition, by comparing it with the sales growth rates of exporting firms in 2006-2007 and 2007-2008, we examine how the decline at the time of the crisis differed from growth rates in normal periods. Table 1 and **Figure 1**(a) show that the distribution of growth rates shifts clearly to the left during the Financial Crisis. In other words, the growth rate at the time of the shock deteriorated substantially relative to normal periods, confirming that this was a large negative shock to exporting firms.

Variable	Mean	SD	P25	Median	P75	Min	Max
2008-2010	-0.1596	0.2576	-0.2913	-0.1189	0.0004	-0.8000	0.4000
2007-2008	0.0357	0.2416	-0.0426	0.0152	0.0892	-4.5197	9.7672
2006-2007	0.0514	0.2913	-0.0280	0.0286	0.1077	-9.5307	7.1399

Table 1: Summary statistics of sales growth rates of exporting firms. The table reports summary statistics of the sales growth rates of exporting firms in 2006-2007, 2007-2008, and 2008-2010.

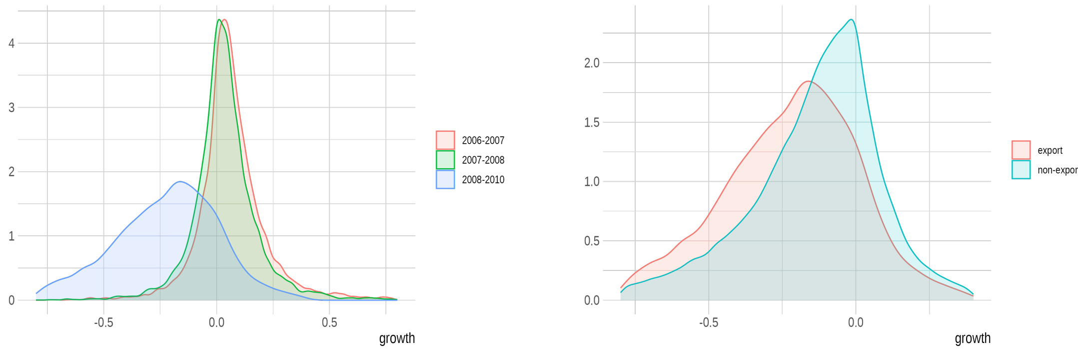

(a) Exporting firms

(b) Exporting and non-exporting firms

Figure 1: Kernel density estimates of growth rates. Panel (a) shows the kernel density estimates of the sales growth rates of exporting firms in 2006-2007, 2007-2008, and 2008-2010. Panel (b) shows the kernel density estimates of the sales growth rates of exporting firms and non-exporting firms in 2008-2010.

Furthermore, when we compare the growth rates of exporting firms and non-exporting firms in 2008-2010, the mean growth rate is -0.2406 for exporting firms and -0.1499 for non-exporting firms. This indicates that the Financial Crisis was a large negative shock to exporting firms. In the next section, we treat this large negative shock to exporting firms as an exogenous shock and analyze, using a GNN, how it propagated to non-exporting firms.

Group	N	mean 2008-2010	mean 2007-2008	mean 2006-2007
Export	7,885	-0.2405	0.0494	0.0755
Non-export	66,196	-0.1500	0.0341	0.0486

Table 2: Sales growth rates of exporting and non-exporting firms. The table reports the mean sales growth rates of exporting firms and non-exporting firms in 2006-2007, 2007-2008, and 2008-2010.

Next, we provide statistics on transaction links. **Figure 2**(a) is a histogram of the degree of links for the non-exporting firms used as prediction targets, including both customers and suppliers. The average degree is 17.17.

Figure 2(b) is a histogram of the share of exporting firms among each firm's transaction partners, that is, exposure. As shown in the figure, the most common case is exposure equal to 0, meaning that none of the firm's transaction partners are exporting firms. In fact, the

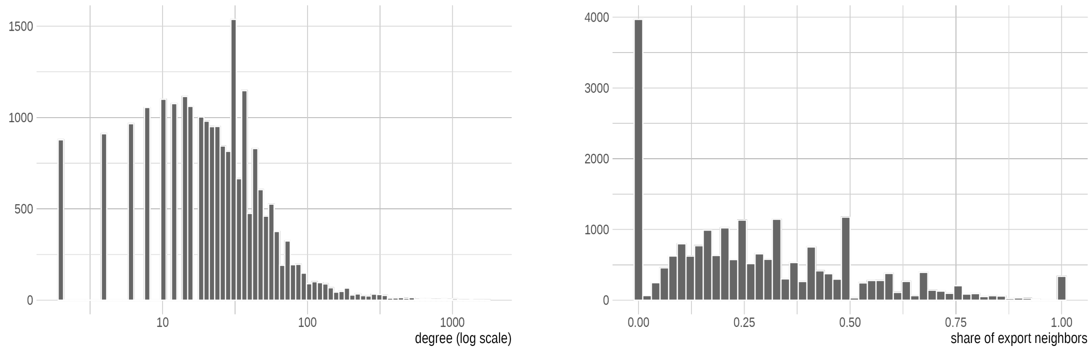

(a) Network degree **(b)** Exposure

Figure 2: Histograms of network degree and exposure Panel (a) shows the histogram of the degree of links for the non-exporting firms used as prediction targets. Here, degree is defined as the sum of the number of suppliers and customers of each firm. Panel (b) shows the histogram of the share of exporting firms among the transaction partners of each non-exporting firm used as a prediction target.

average exposure is 0.281. Therefore, our analysis covers not only cases in which shocks are directly transmitted through transaction partners that are exporting firms, but also cases of indirect shock propagation.

4 Results

We evaluate the predictive accuracy of the GAT model trained using the training data described in Section 3.2. We measure the model's accuracy using the following four indicators.

$$\text{Mean Squared Error (MSE)} = \frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2,$$
$$\text{Mean Absolute Error (MAE)} = \frac{1}{n}\sum_{i=1}^{n}\left|y_i - \hat{y}_i\right|$$
$$\text{Pearson's Correlation r} = \frac{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)\left(\hat{y}_i - \bar{\hat{y}}\right)}{\sqrt{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)^2}\sqrt{\sum_{i=1}^{n}\left(\hat{y}_i - \bar{\hat{y}}\right)^2}}$$
$$R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)^2}$$

Since the target growth rate for 2008-2010 is standardized to have a standard deviation of 1, if the model had no explanatory power at all, then MSE $= 1$ and MAE ≈ 0.8.

The results of the trained model on the test data are MSE $= 0.617$, MAE $= 0.563$, $r = 0.54$, and $R^2 = 0.286$. **Figure 3** plots the true values and the predicted values from the model on the test data. These results indicate that the model has non-trivial explanatory power for firm growth rates. In general, predicting firm growth rates is difficult, and especially under a large-scale shock such as the Financial Crisis, the magnitude of the shock's impact can differ substantially even among firms on the same network. Nevertheless, the fact that the model shows a certain degree of explanatory power on the test data suggests that part of the variation in firm growth rates is predictable using network information.

Next, in order to show that the explanatory power of the trained model above comes from trading partner information, that is, from the propagation of shocks from exporting firms, we train a model with all transaction links masked and examine its predictive accuracy.[4] Keeping the settings and learning method the same, the results on the test data for the model trained with all transaction links masked are MSE $= 0.680$, MAE $= 0.594$, $r = 0.471$, and $R^2 = 0.213$. The predictive accuracy of the model with masked transaction links deteriorates substantially relative to the model that incorporates transaction links. This difference in predictive accuracy indicates that transaction links, namely information on trading partners, are important for prediction. In particular, R^2 increases by about 7.3%, suggesting that part of the predictability of growth rates comes from the transaction network.

Finally, using the model trained with transaction links included, we examine what kinds

[4]Text data about firms is also excluded here because it contains information about their business partners.

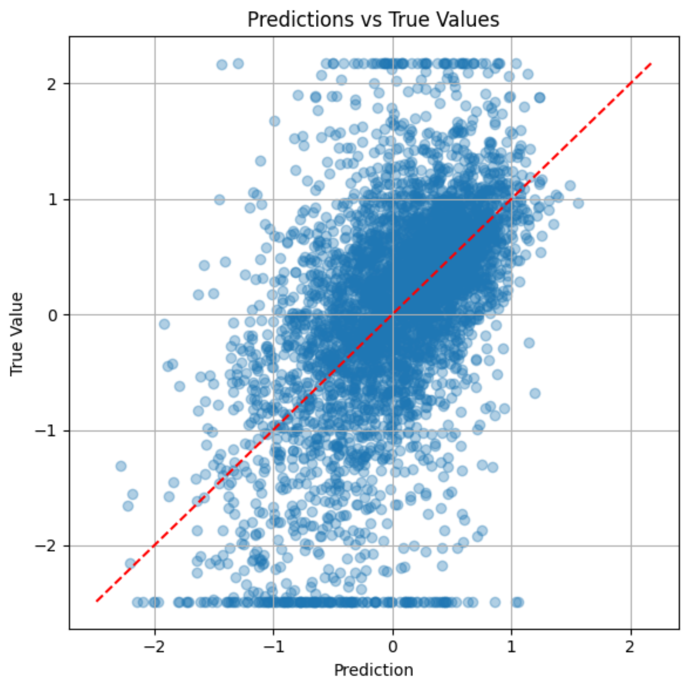

Figure 3: Plot of the model's predicted values against the true values. The figure uses test data that were not used for training. The horizontal axis shows the predictions from the trained model, and the vertical axis shows the true values. For reference, a 45-degree line is also drawn.

of trading partners are identified as important. Specifically, we measure link importance by how much the prediction changes when a link is removed. Here, focusing on a steel company, we extract the top eight transaction relationships that are most important for prediction and examine their industry classifications as well as whether they are customers or suppliers.

The results are reported in **Table 3**. Seven of the top eight firms are customers. This suggests that the trained model regards shocks from the demand side, rather than the supply side, as important for predicting growth rates. This result is consistent with demand-shock propagation. In addition, as shown in the table, firms closely related to the steel company, such as iron and steel wholesalers and related firms, appear at the top. Although this is only one example, it is consistent with the interpretation of demand-shock propagation.

Importance	Type	Industry
0.8486	Customer	Petroleum wholesale, mineral wholesale, and other chemical products wholesale
0.8480	Customer	Steelmaking and steel rolling
0.8426	Customer	Iron and steel wholesale, petroleum wholesale, and gasoline stations
0.8426	Customer	Integrated circuit manufacturing, semiconductor device manufacturing, and electron tube manufacturing
0.8419	Supplier	Chemical machinery manufacturing, industrial furnace manufacturing, and other general industrial machinery and equipment manufacturing
0.8416	Customer	Iron and steel wholesale and hot rolling
0.8413	Customer	Iron and steel wholesale and non-ferrous metal wholesale

Table 3: Top eight important links. For one steel company, the table lists the top eight trading partners that are important for its growth rate, ranked in descending order of importance. Importance takes a value between 0 and 1, with values closer to 1 indicating greater importance.

5 Conclusion

This study presents an empirical framework for predicting shock propagation in interfirm transaction networks by using GNNs, in particular a GAT. Although conventional regression analysis and standard machine learning can make predictions based on firm attributes, they do not easily allow us to explicitly address through which trading partners shocks are strongly transmitted, or which firms and links are important as propagation paths. By contrast, the framework used in this study incorporates the transaction network structure itself into the model in addition to firm attributes, and, by using attention, to simultaneously provide prediction and interpretation that reflect heterogeneity in interfirm relationships.

As a case study, the sharp decline in exports during the financial crisis is treated as an exogenous demand shock, and the analysis examines how this shock propagated to domestic non-exporting firms in Japan. The empirical results confirm that the trained GAT model has a certain degree of explanatory power for predicting firm growth rates on the test data. Moreover, a substantial improvement in predictive performance is confirmed relative to a model in which all transaction links are masked. This indicates that the transaction network structure contains additional information for predicting sales growth rates. Although firm growth rates are inherently difficult to predict, obtaining this level of accuracy out of sample suggests that part of shock propagation is predictable through information on the transaction network. In addition, when important links are extracted through an analysis of attention, most of the top important trading partners for one steel firm are customers rather than suppliers. This result is consistent with the economic interpretation that the main shock in Japan during the financial crisis was on the demand side, and it suggests that the important links captured inside the trained model have not only statistical meaning but also economic implications.

Several issues remain for future research. First, there is still room to improve predictive accuracy, and higher performance will require richer input variables, revisions to the model structure, and further tuning of hyperparameters. Second, although importance measures based on attention and link removal are useful as interpretive tools, additional analysis is necessary before giving them a strict causal interpretation. Third, while this study focuses on the sharp decline in exports during the financial crisis, it is necessary to examine whether the same framework is also effective for supply shocks such as the Great East Japan

Earthquake and for shocks like the COVID-19 pandemic that combine both demand-side and supply-side elements. This study is only an initial attempt, but by extending the analysis to other shocks and refining the interpretation methods, the approach is expected to develop into a practical simulation tool that can contribute to policymaking and risk management.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A. E., and Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016.

Baqaee, D. R. and Farhi, E. (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica*, 87(4):1155–1203.

Barrot, J. and Sauvagnat, J. (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics*, 131(3):1543–1592.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Brody, S., Alon, U., and Yahav, E. (2022). How attentive are graph attention networks? In *10th International Conference on Learning Representations, ICLR 2022*.

Carvalho, V. M. (2014). From Micro to Macro via Production. *Journal of Economic Perspectives*, 28(4):23–47.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east Japan earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321.

Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production Networks: A Primer. *Annual Review of Economics*, 11:635–663.

Fujiy, B. C., Ghose, D., and Khanna, G. (2024). Production networks and firm-level elasticities of substitution. *Working Paper*.

Herskovic, B., Kelly, B., Lustig, H., and Nieuwerburgh, S. V. (2020). Firm volatility in granular networks. *Journal of Political Economy*.

Horvath, M. (1998). Cyclicality and sectoral linkages: Aggregate fluctuations from independent sectoral shocks. *Review of Economic Dynamics*, 1(4):781–808.

Horvath, M. (2000). Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics*, 45(1):69–106.

Huneeus, F. (2020). Production network dynamics and the propagation of shocks. *Working Paper*.

Long, J. B. and Plosser, C. I. (1983). Real business cycles. *Journal of political Economy*, 91(1):39–69.

Peter, A., Ruane, C., et al. (2023). The aggregate importance of intermediate input substitutability. *Working Paper*.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. (2018). Graph attention networks. *arXiv preprint arXiv:1710.10903*.

Wasi, A. T., Islam, M., Akib, A. R., and Bappy, M. M. (2024). Graph neural networks in supply chain analytics and optimization: Concepts, perspectives, dataset and benchmarks. *arXiv preprint arXiv:2411.08550*.

Yang, S., Ogawa, Y., Ikeuchi, K., Shibasaki, R., and Okuma, Y. (2024). Post-hazard supply chain disruption: Predicting firm-level sales using graph neural network. *International Journal of Disaster Risk Reduction*, 110:104664.