B2G Auctions and the Classical Law of Supply and Demand: An XAI-Enabled Deep Learning Approach

Victor Kwan* and Chong Guan

Singapore University of Social Sciences, S599494, Singapore
Email: victorkwanyh@suss.edu.sg (V.K.); guanchong@suss.edu.sg (C.G.)
*Corresponding author

Manuscript received January 25, 2025; accepted May 2, 2025; published September 19, 2025.

***Abstract*—This study explores the rationality of B2G auctions through Singapore's Certificate of Entitlement (COE) system, where car ownership is regulated via online auctions. Utilizing 20 years of data, a XAI-enabled deep learning approach reveals the influence of supply and historical prices on bidding outcomes. These findings align B2G auctions with classical economic theories, which is unlike irrational auction behaviors found in C2C and B2C contexts. This research contributes to auction theory and offers policy makers applicable insights into pricing strategies.**

***Keywords*—auction theory, B2G auctions, deep learning approach, pricing strategies, rational auction behaviors**

I. Introduction

A. Irrational Behavior in Auctions

Auction systems are used in several industries. These include real estate, rare collector's items, commodities, and e-commerce. Some of these biddings are done between Consumer to Consumer (C2C) or Business to Consumer (B2C), while others are Business to Business (B2B) or Business to Government (B2G).

At their most fundamental level, auctions are designed based on economic principles such as tatonnement and game theories (Milgram, 2000) as well as the classical law of supply and demand (we will refer to it as the Law). In its simplest form, the Law states that the price of any commodity will be determined by the level of supply and demand for it. If demand exceeds supply, the price will rise to slow demand down to a level that matches supply. And if supply exceeds demand, the price will drop, leading to increasing demand and clearing excess inventory (Gale, 1955). Auction is arguably the ideal platform to exhibit the workings of this Law.

In recent times, many modern economists and researchers have challenged this classical law of supply and demand. Buechner (2018) explains why the Law has been largely disregarded in graduate economics and proposes that prices are determined differently in different market structures. Similarly, recent research has also discovered that actual auction behavior does not adhere to the law of supply and demand. Factors such as human emotions, hedonic motivations, and risk adversity have significant impacts on the bidding process and results (Budish and Takeyama, 2001; Chang and Chen, 2015; Raviv, 2006). As discussed in contemporary pricing theories, historical bid prices are also unreliable indicators of current ending prices (Liu and Leszczyc, 2023). These economically irrational behaviors are especially pervasive when consumers are involved in the auction, that is, in C2C and B2C bid settings. In fact, empirical evidence has also shown that such abnormalities are even more prevalent in C2C scenarios (Oh, 2014).

This raises the question of whether such economically irrational behaviors are similarly played out in B2B and B2G auction contexts. Businesses are typically more structured, data-driven, and less emotional in their decision-making. As such, the classical law of supply and demand as well as pricing theory should apply more consistently when businesses are bidding parties. To date, there are limited studies on whether irrationality influences B2B and / or B2G auctions or whether do principles of the Law and classical pricing theory hold true. As auction is an important pricing mechanism in many high-value B2B and B2G transactions such as commodity trading and government land sales to developers, there are important practice and research implications to this inquiry. This, therefore, is the aim of our study and we have three hypotheses.

H1: Supply has a significant influence on bid prices when businesses are parties to an auction process.

H2: Demand has a significant influence on bid prices when businesses are parties to an auction process.

H3: Historical prices have a significant influence on bid prices when businesses are parties to an auction process.

B. Car Ownership in Singapore

To conduct our study, we investigate the unique car ownership system in Singapore. Due to the scarcity of space, Singapore's government controls the number of cars on the roads through a Certificate of Entitlement system. Anyone who wants to buy a new car must first secure this certificate (or COE for short) that grants the right to drive it. Each COE is tagged to a car and has a life span of 10 years. The government issues a new piece of COE only when an old one expires and the corresponding car is removed from the roads. Thus, the total car population in the country is controlled with this one-for-one replacement system that does not allow any growth. Allocation of COEs is determined through an online auction system conducted on a bi-weekly basis and bid mostly by car dealers. Thus, making this a B2G or G2B auction context (for simplicity, we will refer to COE bidding as B2G).

Here, the element of supply is the quantity of COEs made available in each bidding exercise and demand is the number of bids placed in the same exercise. As this is an open, online bidding system, supply, demand, and historical ending prices data are transparent and publicly available for every bidding exercise in the last 20 years. With this system, dealers must do forward pricing to maintain their margins. Supply quantity is made known well before each bidding, while they consider factors such as new model launches and buying sentiments to

doi:10.18178/joebm.2025.13.3.885

forecast demand. Past strike prices are also taken into consideration. All these are often done in a structured and data-driven manner. In line with our hypotheses, we expect COE prices to have strong correlations with supply, demand and historical prices, following closely the classical law of supply and demand.

The remainder of this paper is organized as follows: a literature review followed by the methodology of our data collection, analysis, results, conclusions, and limitations.

II. Literature Review

A. *Law of Supply and Demand*

Auction has been defined as a bidding mechanism with a set of rules to determine a winner and the amount he must pay (Wolfstetter, 1996). While there are various types of auctions in practice such as the English auction, Dutch auction, and Vickrey auction (Klemperer, 1999), the fundamentals of auctions are largely based on economic theories including the law of supply and demand (Milgram, 2000). This economic theory of supply and demand is often credited to economist Adam Smith and his seminal 1776 work "The Wealth of Nations".

Over the decades, many scholars in this field have further studied auction systems and integrated other theoretical models to help us understand bidding outcomes that cannot be completely explained by supply and demand alone. These models include the second price auction format (Vickrey, 1961), Nash Equilibrium (Nash, 1951) and Bayesian Game Theory (Harsanyi, 1967). All these early works and more from later years form what we now collectively refer to as the "Auction Theory".

B. *Irrational Behavior of Consumers*

Besides behavioral economics, consumer behavior is another factor that can have a significant influence on bid prices. Chang and Chen (2015) investigate the roles of hedonic and utilitarian motivations in C2C auctions and how time pressure and competition moderate their impact. They find that when under time pressures, bidders process product information more efficiently and this leads to more utilitarian-motivated decisions. On the other hand, consumers tend to be more hedonically driven when faced with bidding competition. The presence of competitors not only adds to the excitement of auctions, it also increases perceived value of the item on auction. Thus, with these psychological behavior, the final bid prices may have little correlation with demand, supply and even past prices.

Winner's curse is a common catchphrase used in auctions. It refers to the phenomenon where a bid winner ends up overpaying for an item due to incomplete information and overly optimistic valuation. Logically, for auctions conducted over digital platforms, this phenomenon should not occur as information on actual, fixed price value is readily available (Mehta and Lee, 1999). However, Oh (2014) finds that a significant number of online bidders still suffer from winner's curse, especially in C2C cases. This can be due to reasons such as overpaying purely for the enjoyment of a competitive game of auction, loss aversion, or perceptual increase in the value of a product based simply on new bid amounts – all of which are irrational behavior that once again deviate from classical economic theories.

In rational settings, historical bid prices should also influence the final prices of subsequent auctions (Dolansky and Vandenbosch, 2013). However, Liu and Leszczyc (2023), building on the foundations of auction and pricing theories, find that when information on a range of historical prices is available, the highest price in the range has an impact on the current price. Moreover, the more information bidders have on historical prices, the higher the ending prices for subsequent bids. Beside these irrational behaviors, guide prices and posted prices can also influence final bid prices even if they are arbitrarily set (Truong *et al.*, 2023; Ryan and Brannigan, 2022).

C. *Limited Research on B2B and B2G Auctions*

Existing literature on auctions primarily focuses on consumers due to practical and theoretical considerations related to market dynamics, data availability, and research priorities. As a result, studies in B2B and B2G contexts remain relatively limited (Dolpanya *et al.*, 2008). This is especially the case for B2G.

There is some research work done in this area and these include Truong *et al.* (2023) who investigate multichannel B2B markets where high posted prices in presales events have a positive influence on subsequent bid prices. This is the case even when no sales are done in the presales phase. Dolpanya *et al.* (2008) look at antecedents to suppliers' participation in B2G auctions, with the hope that their findings will encourage more research for B2G contexts.

III. Methods

A. *Data*

The dataset used in this study was obtained from the Land Transport Authority (LTA) of Singapore, containing bi-weekly observations of key variables in COE bidding: Quota available (supply), bids received (demand), and price results (premium) for each bidding exercise for Categories A (Small cars as defined by LTA) and B (Larger cars as defined by LTA) between 2003-04-01 and 2024-10-14. We focused on Categories A and B in the empirical study as these are the key clusters from which COEs for passenger cars are supplied. We excluded commercial vehicles and motorcycles from our study to reduce unnecessary complexity and they account for a smaller proportion of the total vehicle population.

Other than the three principal variables – premium, quota, and bids received, personal saving rate was extracted from the Singapore Department of Statistics (Singapore Department of Statistics, 2024). Due to its relevance to overall economic health, consumer confidence and purchasing power, personal saving rate was included in the model as a covariate. Raw data were aggregated and aligned into a unified time series, ensuring consistent timestamps across all variables, resulting in 518 usable observations in time periods.

B. *LSTM Model Architecture and Training*

Each feature was then scaled to the [0, 1] range using the MinMaxScaler (Pedregosa *et al.*, 2011) to facilitate stable neural network training. Given the temporal nature of the data, a sliding window approach was used to transform the time series into supervised learning sequences. Specifically, each input (premium, quota, bids received, and personal saving rate) to the model consisted of a window of $seq_{length} = 6$ past timesteps for each of the four variables, while the target was defined as the $premium$ value at the subsequent (future) timestep. This transformation yielded an input array of shape $(num\_samples, 6,4)$ and a corresponding 1D target array. The transformed dataset was then split into training and testing subsets, using 80% of the data for training and 20% for testing.

A Long Short-Term Memory (LSTM) neural network (Hochreiter, 1997) was employed to capture the temporal dependencies in the data. The specific network configuration was determined through an extensive hyperparameter search (e.g. via Keras Tuner complemented by a manual grid search). The final model comprises of: Input Layer: Accepting sequences of length 6 across 4 features; Bidirectional LSTM Layers: Two stacked Bidirectional LSTM layers (Schuster and Paliwal, 1997), each with between 50–200 hidden units, capturing both forward and backward temporal patterns; Additive attention mechanisms were tested to emphasize critical timesteps within the sequence; Dropout Regularization: Set at 40% to mitigate overfitting; and Fully Connected Output Layer: Producing a single prediction for the premium value one timestep ahead.

During training, the Mean Squared Error (MSE) was used as the loss function, optimized via the Adam optimizer with varying learning rates. Early stopping and ReduceLROnPlateau callbacks were implemented to avoid overfitting and adapt the learning rate dynamically. The model was trained for a maximum of 200 epochs but typically converged earlier.

The trained LSTM was assessed on the held-out test set using Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) as primary performance metrics. These metrics quantify the average and root-mean-square magnitudes of the prediction errors, respectively, offering a comprehensive evaluation of predictive accuracy (See Table 1).

Table 1. Model evaluation

	MAE	RMSE
Category A	3080.78	4339.83
Category B	5634.28	7555.09

As further elaborated in Figs. 1 and 2, the models explain both premiums of Categories A and B relatively well.

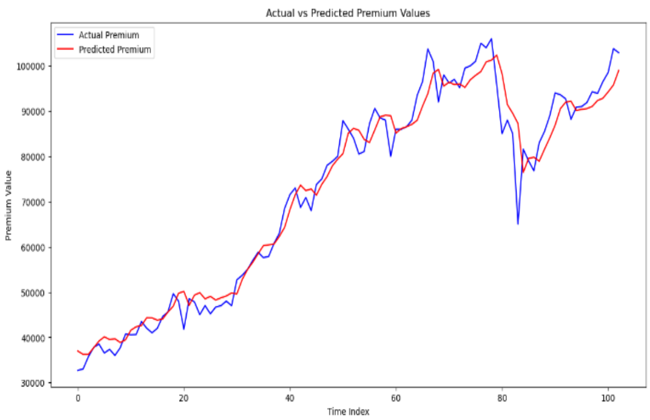

Fig. 1. Actual and predicted premium values for category A.

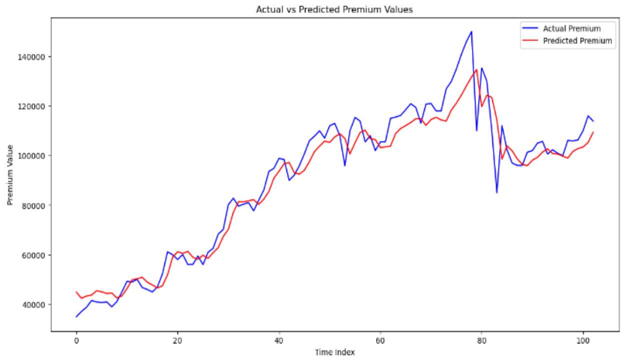

Fig. 2. Actual and predicted premium values for category B.

C. SHAP-Based Model Interpretation

While the LSTM inherently captures temporal dependencies, its internal decision-making process often lacks transparency. To address this limitation, we employed SHAP (SHapley Additive exPlanations) (Lundberg and Lee, 2017) to interpret the feature contributions at each timestep. We leveraged a model-agnostic SHAP explainer—KernelExplainer—thereby ensuring interpretability regardless of the model's architecture.

Table 2. SHAP values for category A

Feature	Mean SHAP Value	t-statistic	p-value
t5_premium	0.15612	14.433	0.000
t4_premium	0.07337	14.175	0.000
t3_premium	0.03358	14.125	0.000
t2_premium	0.01455	13.928	0.000
t1_premium	0.00566	13.302	0.000
t5_quota	0.00394	10.410	0.000
t4_quota	0.00201	9.518	0.000
t0_premium	0.00140	10.360	0.000
t0_personal_saving_rate	0.00098	1.030	0.305
t3_quota	0.00091	8.249	0.000
t2_quota	0.00066	3.177	0.002
t2_bids_received	0.00048	0.740	0.461
t1_quota	0.00019	3.983	0.000
t0_quota	0.00018	2.135	0.035
t5_personal_saving_rate	0.00015	1.756	0.082
t1_personal_saving_rate	0.00007	1.469	0.145
t2_personal_saving_rate	0.00005	1.079	0.283
t0_bids_received	0.00004	1.132	0.260
t3_personal_saving_rate	0.00000	0.126	0.900
t4_personal_saving_rate	−0.00001	−0.232	0.817
t1_bids_received	−0.00005	−1.030	0.305
t3_bids_received	−0.00049	−4.239	0.000
t4_bids_received	−0.00130	−7.414	0.000
t5_bids_received	−0.00302	−8.235	0.000

To accommodate SHAP's requirements for a 1D or 2D input, we reshaped the 3D LSTM input $(samples, 6,4)$ into a 2D representation $(samples, 24)$, then defined a custom prediction function that could reshape it back to 3D when passing data into the LSTM. SHAP values were computed for each of the 24 "flattened" features, which represent timesteps $(t0, t1, \ldots, t5)$ across the four original variables. These values quantify how each feature at each timestep shifts the model's prediction relative to the average prediction. High positive SHAP values indicate strong positive contributions, whereas negative SHAP values indicate a suppressive effect on the predicted premium.

Table 3. SHAP values for category B

Feature	Mean SHAP Value	t-statistic	p-value
t5_premium	0.15671	15.978	0.000
t4_premium	0.07105	15.941	0.000
t3_premium	0.03117	15.431	0.000
t5_quota	0.01387	16.112	0.000
t2_premium	0.01276	14.576	0.000
t4_quota	0.00654	13.671	0.000
t1_premium	0.00458	13.057	0.000
t3_quota	0.00270	12.013	0.000
t2_quota	0.00139	4.966	0.000
t0_premium	0.00126	6.995	0.000
t0_quota	0.00043	1.771	0.080
t0_personal_saving_rate	0.00043	0.895	0.373
t1_quota	0.00035	4.122	0.000
t0_bids_received	0.00003	0.351	0.727
t1_personal_saving_rate	0.00000	0.032	0.975
t5_personal_saving_rate	−0.00002	−0.186	0.853
t3_personal_saving_rate	−0.00005	−0.699	0.486
t2_personal_saving_rate	−0.00006	−0.580	0.563
t4_personal_saving_rate	−0.00009	−0.740	0.461
t2_bids_received	−0.00043	−0.869	0.387
t1_bids_received	−0.00047	−2.049	0.043
t3_bids_received	−0.00283	−10.701	0.000
t4_bids_received	−0.00679	−13.290	0.000
t5_bids_received	−0.01484	−15.255	0.000

After obtaining SHAP values for each feature-timestep combination, we performed one-sample t-tests against zero to determine whether their mean contribution was statistically significant (see Tables 2 & 3). Low p-values indicate that the average impact of a given feature at a specific timestep was unlikely to be due to chance, reinforcing its importance in the LSTM's decision-making process.

D. Results

Across both Category A (Table 2) and Category B (Table III), past premium consistently exhibits the highest positive SHAP values and extremely low p-values, indicating that past premium levels are the strongest predictor of the model's forecasts for future premium. This, therefore, provides evidence to support H3. For quota, moderate positive SHAP values also achieve high significance at earlier timesteps (e.g., t5_quota), underscoring a secondary but nontrivial influence on the predicted premium. As such, there is also statistical evidence to support H1. By contrast, bids received predominantly display negative SHAP contributions—particularly at t5—suggesting that higher bids received in earlier periods may be inversely related to subsequent premium values. Hence, H2 is not supported. Meanwhile, the personal saving rate consistently shows near-zero mean SHAP values and statistically non-significant p-values, indicating a negligible effect on premium forecasts.

Overall, these findings highlight that historical premium remains the dominant predictor, with quota exerting a meaningful but smaller positive effect and bids received having a more modest and negative influence. Beyond historical pricing and the law of supply and demand, personal saving rates does not seem to have a significant impact on bidding results.

IV. Conclusions and Implications

Auctions are designed based on classical economic theories, including the law of supply and demand as well as parts of pricing theory. However, modern scholars have uncovered irrational behaviors that deviate from these classical theories and they have a substantial impact on bidding results. This is especially true when consumers are involved as bidding parties in C2C and B2C contexts.

We hypothesize that since businesses tend to be structured, data-driven, and less emotional in their decision-making, B2B and B2G auction results will be more influenced by supply, demand, historical prices and less so by other factors. With data analysis on a B2G auction scenario, we find support that supply and past prices do have significant influence and, thus, predictive qualities while an external factor such as saving rate has little influence. This shows that B2G auctions are indeed more consistently aligned with classical economic theories.

Contrary to our expectations, we do not find support that demand influences end results. As discussed earlier, dealers do forward pricing due to this auction system and COE prices have a direct impact on their margins. As a result, some of the bids received may be opportunistically put in at low values to try to improve margins. While we do not find support for H2, this is likely due to the profit maximization behavior of businesses and not irrationality.

Overall, our results contribute to auction theory by filling the knowledge gap in B2G contexts. In doing so, we have also added to research literature on behavioral factors affecting auction systems. The application of XAI-enabled deep learning approach in our study also contributes to the small but growing body of research on auction using AI and deep learning. In application, our findings can help governments in pricing strategies and policy making, specifically when an auction is used or is considered as a potential mechanism.

V. Limitations and Future Research

While this study provides valuable insights into auctions, there are inherent limitations that should be considered.

Singapore's COE system is a unique car ownership policy that utilizes auction as an allocation mechanism. Further research on other B2G auctions can be conducted to enhance the generalizability of our findings.

We included only one other covariate that is not related to supply, demand, and historical prices. Future research can include other covariates to further test the rationality of B2G auctions. Lastly, we have been using our findings from B2G data to loosely assume that similar rational behavior can be observed in a B2B context. Again, future research can be done specifically with B2B data.

Conflict of Interest

The authors declare no conflict of interest.

References

Budish, E. B., & Takeyama, L. N. 2001. Buy prices in online auctions: irrationality on the internet? ***Economics Letter***, 72: 325–333.

Buechner, M. N. 2018. A comment on the law of supply and demand. ***The Journal of Philosophical Economics: Reflections on Economics and Social Issues***, XI:2: 67–80.

Chang, C. C., & Chen, C. W. 2015. Examining hedonic and utilitarian bidding motivations in online auctions. ***International Journal of Electronic Commerce***, 19(2):39–65.

Dolansky, E., & Vandenbosch, M. 2013. Price sequences, perceived variability, and choice. ***Journal of Product and Brand Management***, 22(4): 314–321.

Dolpanya, K., L. P. W., & Dick, G. 2008. Antecedents of suppliers' participation in Business-to-Governement (B2G) electronic auction market: Thai B2G E-auction. ***GlobDev 2008.10.***

Gale, D. 2018. The law of supply and demand. ***Maths Scand,*** 3: 155–169.

Harsanyi, J. C. 1967. Games with incomplete information played by "Bayesian" players, I-III. Part 1. The basic model. ***Management Science***, 14(3): 159–1982.

Hochreiter, S. 1997. ***Long Short-term Memory***, Neural Computation MIT-Press, 1997.

Klemperer, P. 1999. Auction theory: A guide to the literature. ***Journal of Economic Surveys***, 13(3): 227–286.

Land Transport Authority (LTA). COE Bidding Results / Prices. https://data.gov.sg/datasets/d_69b3380ad7e51aff3a7dcc84eba52b8a/view (accessed).

Liu, X. T., & Leszczyc, P. T. L. P. 2023. The reference price effect of historical price lists in online auctions. ***Journal of Retailing and Consumer Services***.

Lundberg, S. & Lee, S. I. 2017. SHAP: A unified approach to interpreting model predictions. ***Advances in Neural Information Processing Systems***, 1–10.

Mehta, K., & Lee, B. 1999. An empirical evidence of winner's curse in electronic auctions. ***ICIS 1999 Proceedings***.

Milgram, P. 2000. Putting auction theory to work: The simultaneous acending auction. ***Journal of Political Economy***, 108(2).

Nash, J. 1951. Non-cooperative games. ***Annals of Mathematics***, 54(2): 286–295.

Oh, W. 2014. C2C versus B2B: A comparison of the winner's curse in two types of electronic auctions. ***International Journal of Electronic Commerce***, 6(4): 115–138.

Pedregosa, F. *et al.* 2011. Scikit-learn: Machine learning in Python. ***The Journal of Machine Learning Research***, 12: 2825–2830.

Raviv, Y. 2006. New evidence on price anomalies in sequential auctions: Used cars in New Jersey. ***Journal of Business and Economic Statistics***, 24(3): 301–312.

Ryan, P., & Brannigan, C. 2022. Auction competitive dynamics and Guide (List) prices in a bubble market. ***Accounting, Finance and Governance Review***, 28.

Schuster, M., & Paliwal, K. K. 1997. Bidirectional recurrent neural networks. ***IEEE Transactions on Signal Processing***, 45(11): 2673–2681.

Singapore Department of Statistics. Personal Disposable Income and Saving Third Quarter 2024. Available: https://www.singstat.gov.sg/pdips (accessed.

Truong, M., Gupta, A., Ketter, W., & Heck, E. van. 2023. The effect of posted prices on auction prices: An empirical investigation of a multichannel B2B market. ***MIS Quarterly***, 47(4): 1557–1584.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. ***The Journal of Finance***, 16(1): 8–37.

Wolfstetter, E. 1996. Auctions: An introduction. ***Journal of Economic Surveys***, 10(4).