Online Regression with Instrumental Variables

Regrets under Endogeneity and Bandit Feedback

Debabrota Basu

Équipe Scool, Inria, Université de Lille, CNRS, France

HyAIAI Evaluation Meeting, June 2023

The Trajectory

1. Warm-up: Online Linear Regression
2. Instrumental Variable Regression
3. Linear Bandits under Endoegenity
4. Future Roadmap: Challenges and Opportunities

Learning Price-Consumption Dynamics

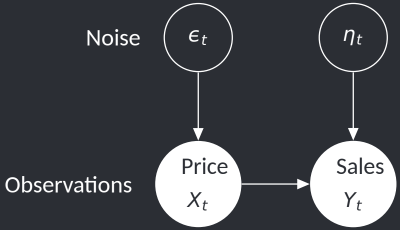

Let us consider

$$\mathsf{Sales}_t = \beta^* \times \mathsf{Price}_t + \eta_t$$

Goal

Learn β^* from a stream of data $\{\mathsf{Sales}_1, \mathsf{Sales}_2, \ldots\}$ & $\{\mathsf{Price}_1, \mathsf{Price}_2, \ldots\}$.

Learning Price-Consumption Dynamics

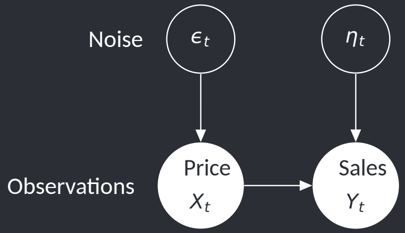

Let us consider

$$\text{Sales}_t = \beta^* \times \text{Price}_t + \eta_t$$

Goal

Learn β^* from a stream of data $\{\text{Sales}_1, \text{Sales}_2, \ldots\}$ & $\{\text{Price}_1, \text{Price}_2, \ldots\}$.

Solution: Online Linear Regression [Wasserman, 2004]

Find the β minimising the square loss till time t

$$\beta_t \triangleq \operatorname*{argmin}_{\beta} \sum_{s=1}^{t} (\text{Price}_s - \beta \times \text{Sales}_s)^2 \,.$$

Online Linear Regression: Premises and Conclusion

Online Linear Regression yields an estimate

$$\beta_t = \left(\sum_{s=1}^{t} \mathsf{Sales}_s\right)^{-2} \times \left(\sum_{s=1}^{t} \mathsf{Sales}_s \times \mathsf{Price}_s\right).$$

We note that β_t is an **unbiased** and **consistent** estimator of β.

Online Linear Regression yields unbiased estimate if

1. The observational noise η_t is independent of Price_t.
2. There is no external or unobserved variable except Price that impacts Sales.

Are these premises always true?

Learning Price-Consumption Dynamics

A Case for Endogeneity

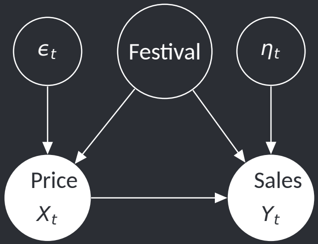

Now, the underlying dynamics is

$$\text{Sales}_t = \beta^* \times \text{Price}_t + \rho_S \times \text{Festival}_t + \eta_t$$

Goal

Learn β^* from a stream of data $\{\text{Sales}_1, \text{Sales}_2, \ldots\}$ & $\{\text{Price}_1, \text{Price}_2, \ldots\}$.

But, online linear regression **does not yield an unbiased estimate that converges to β.** [Wald, 1940, Greene, 2003]

Why Do We Care for Endoegeneity?

Endogeneity is a widely studied phenomenon in epidemiology, economics, bioinformatics, social sciences, and causal inference that emerges due to

- Omitted explanatory variables
  - Estimate the number of returning students to college using the National Survey of Youth data [Rubin, 1974, Carneiro et al., 2011, Mogstad et al., 2021]
- Strategic behaviours during data generation
  - Just-In-Time Adaptive Interventions (JITAI) using mobile health applications [Tewari and Murphy, 2017, Kallus, 2018] (Susan Murphy's plenary talk, AAAI 2023)
- Measurement errors
  - Effect of family income on children's cognitive outcome [Dahl and Lochner, 2012, Zhu et al., 2022]
- Dependence of the output and the covariates on unobserved confounding variables
  - Causal inference with Rubin's potential outcome framework [Rubin, 1974, Angrist and Imbens, 1995, Hernan and Robins, 2020] (Nobel in Econ. 2022)

Tackling Endoegenity in Regression

Introduce Instrumental Variables $\perp\!\!\!\perp$ Unobserved Variables and Noise

First, introduced by Doctor John Snow during **London cholera epidemic of 1853-54** to prove whether **cholera is waterborne**.

Learning Price-Consumption Dynamics under Endoegenity

Introducing Instumental Variables (IVs)

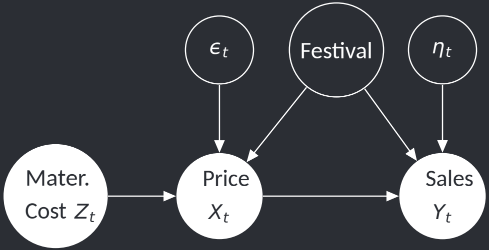

Now, the underlying dynamics has two stages

First stage

$$\text{Price}_t = \theta^* \times \text{MCost}_t + \rho_F \times \text{Festival}_t + \epsilon_t$$

Second stage

$$\text{Sales}_t = \beta^* \times \text{Price}_t + \rho_S \times \text{Festival}_t + \eta_t$$

Goal

Learn β^* and θ^* from a stream of data $\{\text{Sales}_1, \text{Sales}_2, \ldots\}$, $\{\text{Price}_1, \text{Price}_2, \ldots\}$, and $\{\text{MCost}_1, \text{MCost}_2, \ldots\}$.

IVs: Premises and Conclusions

IVs should satisfy

- IVs are **exogeneous** w.r.t. both the first and second-stage noise.
- IVs are **relevant** to estimate the first-stage variable (e.g. Material Cost has enough influence on the Price). Mathematically, covariances of IVs and first-stage variables are always non-zero.

IVs: Premises and Conclusions

IVs should satisfy

- IVs are **exogeneous** w.r.t. both the first and second-stage noise.
- IVs are **relevant** to estimate the first-stage variable (e.g. Material Cost has enough influence on the Price). Mathematically, covariances of IVs and first-stage variables are always non-zero.

IVs lead to

1. an unbiased estimate θ_t of θ^* as in classic online linear regression
2. a predictive value of first-stage variable $\widehat{\mathsf{Price}}_t = \theta_t \times \widehat{\mathsf{MCost}}_t$
3. a decoupling of the second-stage noise η_t and second-stage variable Sales_t given the prediction $\widehat{\mathsf{Price}}_t$
4. an unbiased estimate β_t of β^* through another online linear regression

Online Instrumental Variable Regression

Online Two-stage Least Squares (O2SLS) [Vecchia and Basu, 2023]

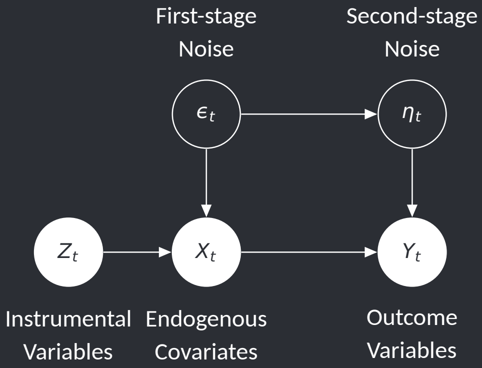

First stage

$$\text{Price}_t = \theta^* \times \text{MCost}_t + \epsilon_t$$

Regression 1: Learn θ^* from a stream of IVs Z_t and covriates X_t

Second stage

$$\text{Sales}_t = \beta^* \times \text{Price}_t + \eta_t$$

Regression 2: Learn β^* from a stream of outcomes Y_t and predicted covariates from the first-stage $\hat{X}_t$.

Online Instrumental Variable Regression

Online Two-stage Least Squares (O2SLS) [Vecchia and Basu, 2023]

Algorithm O2SLS

Input: Initialisation parameters β_0, θ_0
for $t = 1, 2, \ldots, T$ **do**
    Observe z_t generated i.i.d. by Nature, and x_t sampled for given z_t
    Compute first-stage parameter estimates $\theta_{t-1} = (\sum_{s=1}^{t-1} z_s z_s^\top)^{-1} \sum_{s=1}^{t-1} z_s x_s^\top$
    Use θ_{t-1} and z_t to predict $\hat{x}_t$
    Compute second-stage parameter estimates $\beta_{t-1} = (\sum_{s=1}^{t-1} \hat{x}_s \hat{x}_s^\top)^{-1} \sum_{s=1}^{t-1} \hat{x}_s^\top y_s$
    Predict $\widehat{y}_t = \beta_{t-1}^\top x_t$
    Observe y_t generated by Nature
end for

Confidence of Estimating β^*

Lemma (Confidence ellipsoid for the second-stage parameters)

For σ_η-sub-Gaussian first stage noise η_t and for all $t > 0$, the true parameter β belongs to the confidence set around the estimator

$$\mathcal{E}_t \triangleq \left\{\beta \in \mathbb{R}^{d_x} : \|\beta_t - \beta\|_{\widehat{H}_t} \leq \sqrt{\mathfrak{b}_t(\delta)}\right\}, \tag{1}$$

with probability at least $1-\delta \in (0,1)$. Here, $\mathfrak{b}_t(\delta) \triangleq \frac{d_Z \sigma_\eta^2}{4} \log\left(\frac{1+tL_z^2/\lambda d_Z}{\delta}\right)$.

Thus, the confidence ellipsoid around the estimator contracts at a rate $\mathcal{O}\left(\sqrt{\frac{\log t}{t}}\right)$.

Identification Regret of O2SLS

Identification Regret: The cost of identifying the true parameter β^* is given by

$$\widetilde{R}_T(\beta^*) \triangleq \sum_{t=1}^{T} (x_t^\top \beta_{t-1} - x_t^\top \beta^*)^2.$$

Theorem (Identification regret of O2SLS)

The identification regret of O2SLS satisfies with probability at least $1-\delta$

$$\widetilde{R}_T \leq \sum_{t=1}^{T} \underbrace{\|\beta_t - \beta\|_{\widehat{H}_t}^2}_{\text{Estimation}} \times \underbrace{\|x_t\|_{\widehat{H}_t^{-1}}^2}_{\substack{\text{Second-stage}\\ \text{feature norm}}} \leq \underbrace{\mathfrak{b}_{T-1}(\delta)}_{\mathcal{O}(d_z \log T)} \times \underbrace{\sum_{t=1}^{T} \|x_t\|_{\widehat{H}_t^{-1}}^2}_{\mathcal{O}(d_x \log T)} = \mathcal{O}\left(d_x d_z \log^2(T)\right).$$

Oracle (Predictive) Regret of O2SLS

Oracle Regret: The regret in terms of the quality of prediction is defined as

$$\overline{R}_T(\beta) \triangleq \sum_{t=1}^{T}(y_t - x_t^\top \beta_{t-1})^2 - \sum_{t=1}^{T}(y_t - x_t^\top \beta)^2.$$

Theorem (Oracle regret of O2SLS)

Oracle Regret of O2SLS at step $T > 1$ is upper bounded by (ignoring $\log\log$ terms)

$$\underbrace{\widetilde{R}_T}_{\substack{\text{Identif.}\\ \text{Regret}\\ \mathcal{O}(d_x d_z \log^2 T)}} + \underbrace{\sqrt{\mathfrak{b}_{T-1}(\delta)}}_{\substack{\text{Estimation}\\ \mathcal{O}(\sqrt{d_z \log T})}} \Big(\underbrace{C_1\sqrt{f(T)}}_{\substack{\text{First-stage}\\ \text{feature norm}\\ \mathcal{O}(\sqrt{\log T})}} + \underbrace{C_2\sqrt{2 d_x f(T)} + \sqrt{d_x} C_3}_{\substack{\text{Correlated noise}\\ \text{Concentration term}\\ \mathcal{O}(\sqrt{d_x \log T})}} + \underbrace{\gamma C_4 \sqrt{T}}_{\substack{\text{Correlated noise}\\ \text{Bias term}\\ \mathcal{O}(\gamma\sqrt{T})}} \Big) = \mathcal{O}(\gamma\sqrt{T}).$$

with probability at least $1 - \delta \in (0, 1)$. Here, degree of endogeneity $\gamma \triangleq \|\gamma\|_2 = \|\mathbb{E}[\eta_s \epsilon_s]\|_2$.

Experimental Analysis

Part I: Final Regret over Different Degrees of Endogeneity

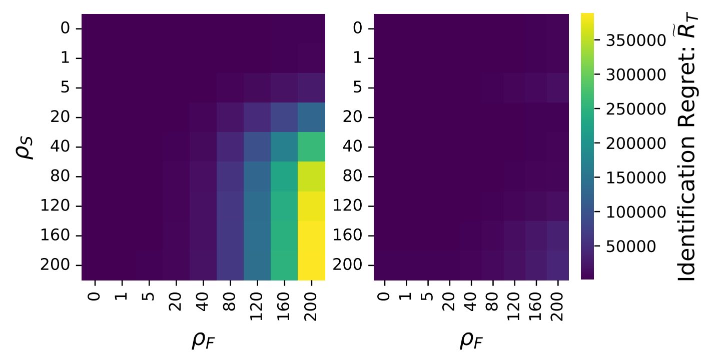

Figure: Identification regret after $T = 10^3$ steps of Online Ridge (left) and O2SLS (right), for different combination of ρ_F and ρ_S in $[0, 200]$. O2SLS attains lower regret than Ridge for a wide range of parameters.

Experimental Analysis

Part II: Evolution of Regret over Different Degrees of Endogeneity

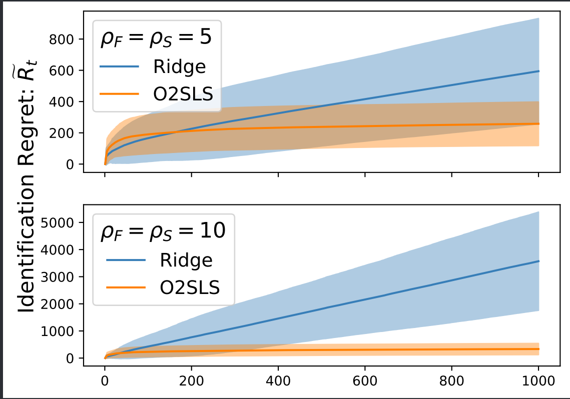

Figure: Identification regret of Online Ridge and O2SLS over $T = 10^3$ steps, and for $\rho_F = \rho_S = 5, 10$. With increase in ρ_S, i.e. endoegenity, O2SLS performs better.

Tackling Endoegenity in Bandits

O2SLS Regression $\Longleftrightarrow$ Optimism in the Face of Uncertainty

Dynamic Pricing under Endoegenity

Bandits with Instrumental Variables [Kallus, 2018, Vecchia and Basu, 2023]

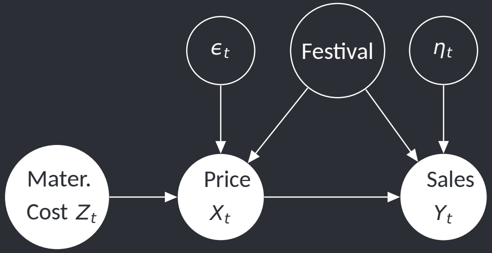

First stage

$$\text{Price}_t = \theta^* \times \text{MCost}_t + \rho_F \times \text{Festival}_t + \epsilon_t$$

Second stage

$$\text{Sales}_t = \beta^* \times \text{Price}_t + \rho_S \times \text{Festival}_t + \eta_t$$

Goal

Given K possible feasible prices between $[0, \text{MaxRetailPrice}]$ and corresponding material costs, selecting which price and which material cost would lead to the highest amount of sales.

Bandits under Endogeneity

Algorithm The Interactive Process of Bandits under Endogeneity

Input: Initialisation parameters $\beta_0, \widehat{\theta}_0$
for $t = 1, 2, \ldots, T$ **do**
    Sample covariates $x_{t,a} \in \mathscr{X}_t$ for all $a \in \mathscr{A}_t$
    Choose an action A_t from the feasible action set $\mathscr{A}_t$
    Observe corresponding IV z_{t,A_t} and outcome y_t
    Update the parameter estimates β_t and θ_t.
end for

OFUL-IV: IVs+Optimism for Bandits under Endogeneity

Algorithm OFUL-IV

for $t = 1, 2, \ldots, T$ **do**

Sample covariates $x_{t,a} \in \mathscr{X}_t$ for all $a \in \mathscr{A}_t$

Compute β_{t-1} using O2SLS estimator

$$\beta_{t-1} \triangleq \left(\sum_{s=1}^{t-1} \widehat{x}_s^\top \widehat{x}_s \right)^{-1} \sum_{s=1}^{t-1} \widehat{x}_s^\top y_s \tag{2}$$

Choose an action A_t from the feasible action set $\mathscr{A}_t$ using optimistic index

$$A_t = \operatorname*{argmax}_{a \in \mathscr{A}_t} \left\{ \langle x_{t,a}, \beta_{t-1} \rangle + \sqrt{\mathfrak{b}'_{t-1}(\delta)}\, \|x_{t,a}\|_{\widehat{H}_{t-1}^{-1}} \right\} \tag{3}$$

Observe corresponding IV z_{t,A_t} and outcome y_t

Update the parameter estimates β_t and θ_t, and confidence interval $\mathfrak{b}'_t(\delta)$

end for

Theoretical Analysis

Theorem (Regret upper bound of OFUL-IV)

Under the same assumptions as that of O2SLS, with probability $1 - \delta$ and for horizon $T > 1$, OFUl-IV incurs regret

$$R_T \leq 2\sqrt{T} \underbrace{\sqrt{\mathfrak{b}_{T-1}(\delta)}}_{\text{Estimation } \mathcal{O}(\sqrt{d_z}\log T)} \underbrace{\left(\sum_{t=1}^{T} \|x_{t,A_t}\|^2_{\widehat{H}_t^{-1}}\right)^{1/2}}_{\text{Second-stage feature norm } \mathcal{O}(\sqrt{d_x}\log T)} = \mathcal{O}(\sqrt{d_x d_z T}\log T)$$

For $d_x = d_z$, we retrieve a regret bound of same order as that of classic linear bandits without exoegeneity. This shows efficiency of OFUL-IV to eliminate bias due to endogeneity while decision making.

Experimental Analysis

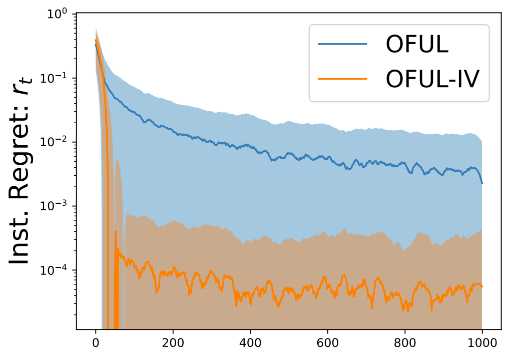

(a) Instantaneous regret

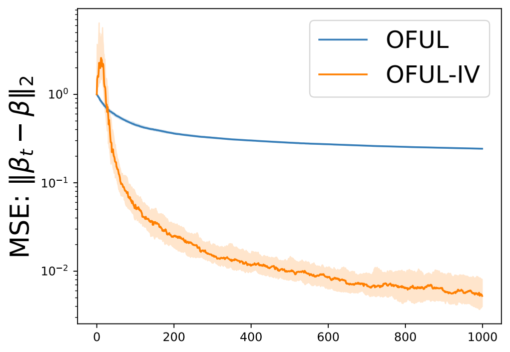

(b) Mean Square Error w.r.t. β

The Road Ahead: Challenges and Opportunities

- **Tackling Non-linearity:** Extending our analysis to non-linear regression problems, like kernel regressions and neural network based regressions
- **Solving Control Problems**: We are working on formulating and solving control problems with underlying causal structures using O2SLS framework as oracle
- **Identifying 'Strong' IVs:** Our analysis depends heavily on existence of a set of strong and relevant IVs. The question is how to identify them or adapt these algorithms when they are weak.

For further details, please visit: `https://debabrota-basu.github.io/`

Questions?

Thanks to Riccardo Della Vecchia,
who has been central to develop these research works.

References I

[Angrist and Imbens, 1995] Angrist, J. D. and Imbens, G. W. (1995).
Two-stage least squares estimation of average causal effects in models with variable treatment intensity.
Journal of the American statistical Association, 90(430):431–442.

[Carneiro et al., 2011] Carneiro, P., Heckman, J. J., and Vytlacil, E. J. (2011).
Estimating marginal returns to education.
American Economic Review, 101(6):2754–81.

[Dahl and Lochner, 2012] Dahl, G. B. and Lochner, L. (2012).
The impact of family income on child achievement: Evidence from the earned income tax credit.
American Economic Review, 102(5):1927–1956.

[Greene, 2003] Greene, W. H. (2003).
Econometric analysis.
Pearson Education India.

[Hernan and Robins, 2020] Hernan, M. and Robins, J. (2020).
Causal Inference: What if.
Boca Raton: Chapman & Hill/CRC.

[Kallus, 2018] Kallus, N. (2018).
Instrument-armed bandits.
In *Algorithmic Learning Theory*, pages 529–546. PMLR.

[Mogstad et al., 2021] Mogstad, M., Torgovitsky, A., and Walters, C. R. (2021).
The causal interpretation of two-stage least squares with multiple instrumental variables.
American Economic Review, 111(11):3663–98.

References II

[Rubin, 1974] Rubin, D. B. (1974).
Estimating causal effects of treatments in randomized and nonrandomized studies.
Journal of educational Psychology, 66(5):688.

[Tewari and Murphy, 2017] Tewari, A. and Murphy, S. A. (2017).
From ads to interventions: Contextual bandits in mobile health.
Mobile Health: Sensors, Analytic Methods, and Applications, pages 495–517.

[Vecchia and Basu, 2023] Vecchia, R. D. and Basu, D. (2023).
Online instrumental variable regression: Regret analysis and bandit feedback.

[Wald, 1940] Wald, A. (1940).
The fitting of straight lines if both variables are subject to error.
The annals of mathematical statistics, 11(3):284–300.

[Wasserman, 2004] Wasserman, L. (2004).
All of statistics: a concise course in statistical inference, volume 26.
Springer.

[Zhu et al., 2022] Zhu, Y., Gultchin, L., Gretton, A., Kusner, M. J., and Silva, R. (2022).
Causal inference with treatment measurement error: a nonparametric instrumental variable approach.
In *Uncertainty in Artificial Intelligence*, pages 2414–2424. PMLR.