Midastar: Threshold autoregression with data sampled at mixed frequencies

Kaiji Motegi[1] John W. Dennis[2]

[1]Kobe University
[2]Institute for Defense Analyses

Econometrics Workshop
Institute for Economic Studies, Keio University
July 5, 2022

Introduction: Background

- A time series often has heterogeneous properties below versus above a certain threshold (**threshold effects**).
- One of the most well-known models in this field is the threshold autoregression (**TAR**) proposed by Tong (1978).
- In TAR, a target variable y follows AR(p) with coefficients being different across regimes, and a regime switch is triggered when a threshold variable x crosses a constant threshold parameter μ.
- Many variants and extensions of TAR have been proposed:
  1. smooth transition autoregression (Granger & Teräsvirta, 1993);
  2. threshold regression (Hansen, 2000);
  3. threshold kink (Hansen, 2017);
  4. time-varying threshold (e.g., Motegi, Cai, Hamori & Xu, 2020; Motegi, Dennis & Hamori, 2022).

Introduction: Motivation

- In the existing threshold models, the target variable y and the threshold variable x are assumed to be sampled at the same frequency.
- In practice, time series are often sampled at different frequencies (e.g., daily, weekly, monthly, quarterly, etc.).
- In the existing literature of threshold models, a variable of the higher frequency is aggregated into the lower level.
- Temporal aggregation has an adverse effect on statistical inference due to the loss of information (e.g., Silvestrini & Veredas, 2008).
- The goal of this project is to extend the TAR model to **mixed frequency data**.

Introduction: Methodology

- To achieve our goal, we rely on the literature of **Mixed Data Sampling (MIDAS)** econometrics originated with Ghysels, Santa-Clara & Valkanov (2004).
- We propose **Midastar** models by combining MIDAS and TAR:
  1. The Midastar model of the first kind (Midastar I):
     low frequency y and high frequency x.
  2. The Midastar model of the second kind (Midastar II):
     high frequency y and low frequency x.
- For both scenarios, x just determines regimes and y is not regressed onto x. This feature makes the Midastar models different from existing MIDAS-type models:
  1. Both Midastar I and Midastar II are easy to formulate.
  2. Parameter proliferation is less an issue.

Introduction: Methodology

- The Midastar models have the regression parameters $\boldsymbol{\beta}$ (i.e., intercepts and AR parameters) and the nuisance parameters γ (i.e., the delay and threshold parameters).
- Partition $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \boldsymbol{\beta}_2^\top)^\top$ by regimes.
- The nuisance parameters γ are unidentified *if and only if* threshold effects are absent (i.e., $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2$).
- When threshold effects are absent, the Midastar models reduce to single-regime AR models.
- When threshold effects are present, the Midastar models are able to detect them, while the low frequency TAR models have a risk of pointing to **spurious non-threshold effects**.

Introduction: Main results

- We estimate (β, γ) via a two-step procedure called **profiling**.
- To test the no-threshold-effect hypothesis, we adopt the wild-bootstrap tests of Hansen (1996).
- The proposed methods have desired statistical properties in large and finite samples.
- We present two separate empirical examples:
    1. Midastar I for Japan's COVID-19 data.
    2. Midastar II for U.S. macroeconomic indicators.
- For both examples, the bootstrap tests based on the Midastar models reject the null hypothesis of no threshold effects, while the tests based on the aggregated TAR fail to reject the null.

Table of contents

Midastar I: Set-up

- Let us begin with a single-frequency framework where time periods are denoted as

$$\mathbb{L} = \{1, \ldots, n\}.$$

- Let $\{y_t\}_{t \in \mathbb{L}}$ be a target variable; let $\{x_t\}_{t \in \mathbb{L}}$ be a threshold variable.

- The two-regime TAR model of Tong (1978) is specified as

$$y_t = \begin{cases} \alpha_1 + \sum_{k=1}^{p} \phi_{1k} y_{t-k} + u_t & \text{if } \; x_{t-d} < \mu, \\ \alpha_2 + \sum_{k=1}^{p} \phi_{2k} y_{t-k} + u_t & \text{if } \; x_{t-d} \geq \mu, \end{cases} \qquad t \in \mathbb{L}.$$

- y has different autocorrelation structures below vs. above the threshold parameter μ.

Midastar I: Motivation

- Now assume the target variable y is observed at a low frequency and the threshold variable x is observed at a high frequency.
- Assume the ratio of sampling frequencies, m, is known and fixed across time. (Example: $m = 3$ if y is sampled quarterly and x is sampled monthly.)
- Define the set of high frequency time periods as $\mathbb{H} = \cup_{t \in \mathbb{L}} \mathbb{H}_t$, where

$$\mathbb{H}_t = \left\{ t - 1 + \frac{1}{m},\ t - 1 + \frac{2}{m},\ \ldots,\ t \right\}, \qquad t \in \mathbb{L}.$$

- For each $t \in \mathbb{L}$, we observe a single realization y_t for the target variable, while we sequentially observe $\{x_j^*\}_{j \in \mathbb{H}_t}$ for the threshold variable.

Midastar I: Motivation

- Let $\{w_1, \ldots, w_m\}$ be the pre-specified linear aggregation scheme such that $w_k \geq 0$ for all k and $\sum_{k=1}^{m} w_k = 1$.
- Two well-known examples of the linear aggregation scheme:
  1. Stock aggregation: $w_k = \mathbf{1}(k = m)$ for $k \in \{1, \ldots, m\}$.
  2. Averaging: $w_k = 1/m$ for $k \in \{1, \ldots, m\}$.
- For $t \in \mathbb{L}$, let $x_t = \sum_{k=1}^{m} w_k x^*_{t-1+k/m}$ be a temporal aggregation of x^*.
- In the existing literature, the single-frequency TAR model is fitted to $\{y_t, x_t\}_{t \in \mathbb{L}}$ even if $\{x^*_t\}_{t \in \mathbb{H}}$ are observable.
- In the MIDAS literature, it is well known that such a temporal aggregation often has an adverse impact on inference due to the loss of information on x^*.

Midastar I: Specification

- To avoid the temporal aggregation of x^*, we propose **the Midastar model of the first kind (Midastar I)**:

$$y_t = \begin{cases} \alpha_1 + \sum_{k=1}^{p} \phi_{1k} y_{t-k} + u_t & \text{if } \; x^*_{t-\frac{d}{m}} < \mu, \\ \alpha_2 + \sum_{k=1}^{p} \phi_{2k} y_{t-k} + u_t & \text{if } \; x^*_{t-\frac{d}{m}} \geq \mu, \end{cases} \quad t \in \mathbb{L}.$$

- The delay of d high frequency periods is taken from the integer time period t, exploiting the high frequency observations of x^*.
- The choice set of the delay parameter d is $\mathcal{D} \subseteq \mathbb{N}$.
- The choice set of the threshold parameter μ is $\mathcal{X}^* = \{x_t^*\}_{t \in \mathbb{H}}$.

Midastar I: Threshold effects

- Stack the regression parameters for each regime:

$$\boldsymbol{\beta}_r = (\alpha_r, \phi_{r1}, \ldots, \phi_{rp})^\top, \qquad r \in \{1, 2\}.$$

- Threshold effects are **absent** if $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2$.
- Threshold effects are **present** if $\boldsymbol{\beta}_1 \neq \boldsymbol{\beta}_2$.
- If threshold effects are absent, the high frequency threshold variable x^* is irrelevant and the Midastar I and the aggregated TAR are essentially equivalent. Hence, the two models are equally capable of detecting the truth of no threshold effects.

Midastar I: Threshold effects

- If threshold effects are present, the aggregated TAR model is generally misspecified relative to the Midastar I, since the former cannot capture the high frequency delay.
- The misspecification results in the failure to identify the true value of $\boldsymbol{\beta}$.
- In particular, the aggregated TAR can reach a wrong conclusion that $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2$ (**spurious non-threshold effects**).
- This insight is in line with the well-known fact that temporal aggregation tends to **weaken** nonlinearities in original series (e.g., Granger & Lee, 1999).
- In this sense, the Midastar I captures threshold effects more precisely than the aggregated TAR.

Midastar I: Key features

- The number of parameters in the Midastar I does not depend on m, hence a large value of m would not be an issue.
- Parameter proliferation due to large m is a major issue in many strands of the MIDAS literature:
    1. unrestricted MIDAS regression (e.g., Foroni, Marcellino & Schumacher, 2015);
    2. mixed frequency vector autoregression (e.g., Ghysels, 2016; Ghysels, Hill & Motegi, 2016);
    3. regression-based Granger causality tests (e.g., Ghysels, Hill & Motegi, 2020);
- Given the Midastar I, empirical studies with large m are feasible.
- Furthermore, a time-varying m_t can easily be allowed.

Midastar I: Matrix representation

- Stack the regression and nuisance parameters:

$$\boldsymbol{\beta} = \begin{bmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{bmatrix}, \quad \boldsymbol{\gamma} = \begin{bmatrix} d \\ \mu \end{bmatrix}.$$

- Define binary variables which represent the regime:

$$I_{1t}^*(\mu) = \mathbf{1}\left(x_t^* < \mu\right), \quad I_{2t}^*(\mu) = \mathbf{1}\left(x_t^* \geq \mu\right), \quad t \in \mathbb{H}.$$

- Stack the regressors:

$$\boldsymbol{z}_t = \begin{bmatrix} 1 \\ y_t \\ \vdots \\ y_{t+1-p} \end{bmatrix}, \quad \boldsymbol{Z}_t(\boldsymbol{\gamma}) = \begin{bmatrix} \boldsymbol{z}_{t-1} I_{1,t-\frac{d}{m}}^*(\mu) \\ \boldsymbol{z}_{t-1} I_{2,t-\frac{d}{m}}^*(\mu) \end{bmatrix}, \quad t \in \mathbb{L},$$

- The matrix representation of the Midastar I is given by

$$y_t = \boldsymbol{Z}_t(\boldsymbol{\gamma})^\top \boldsymbol{\beta} + u_t, \quad t \in \mathbb{L}.$$

Profiling estimation of the Midastar I

- To estimate the regression parameter $\boldsymbol{\beta}$ and the nuisance parameter $\boldsymbol{\gamma}$, we adopt a two-step procedure called **profiling**.
- If $\boldsymbol{\gamma}$ were given, then the least squares estimator for $\boldsymbol{\beta}$ would be analytically available:

$$\hat{\boldsymbol{\beta}}(\boldsymbol{\gamma}) = \left\{ \sum_{t \in \mathbb{L}} \boldsymbol{Z}_t(\boldsymbol{\gamma}) \boldsymbol{Z}_t(\boldsymbol{\gamma})^\top \right\}^{-1} \left\{ \sum_{t \in \mathbb{L}} \boldsymbol{Z}_t(\boldsymbol{\gamma}) y_t \right\}.$$

- The profiling estimator for $\boldsymbol{\gamma}$ is given by:

$$\hat{\boldsymbol{\gamma}} = \arg\min_{\boldsymbol{\gamma} \in \Gamma} \sum_{t \in \mathbb{L}} \left\{ y_t - \boldsymbol{Z}_t(\boldsymbol{\gamma})^\top \hat{\boldsymbol{\beta}}(\boldsymbol{\gamma}) \right\}^2.$$

- The profiling estimator for $\boldsymbol{\beta}$ is given by $\hat{\boldsymbol{\beta}} = \boldsymbol{\beta}(\hat{\boldsymbol{\gamma}})$.

Profiling estimation of the Midastar I

- Asymptotic properties of the profiling estimator depends crucially on whether threshold effects are absent or present.
- Threshold effects are absent if $\boldsymbol{\beta}_1 = \boldsymbol{\beta}_2$, in which case $\boldsymbol{\gamma}$ is unidentifiable.
- Threshold effects are present if $\boldsymbol{\beta}_1 \neq \boldsymbol{\beta}_2$, in which case $\boldsymbol{\gamma}$ is identifiable.
- Define the no-threshold-effect hypothesis:

$$H_0^* : \boldsymbol{\beta}_1 = \boldsymbol{\beta}_2 \qquad \text{vs.} \qquad H_1^* : \boldsymbol{\beta}_1 \neq \boldsymbol{\beta}_2.$$

Profiling estimation of the Midastar I

Theorem 1 (Profiling estimator)

Under standard regularity conditions, the following are true:

1. $\sqrt{n}\{\hat{\boldsymbol{\beta}}(\boldsymbol{\gamma}) - \boldsymbol{\beta}_0\} \Rightarrow \mathcal{N}\{\mathbf{0}, \boldsymbol{V}(\boldsymbol{\gamma})\}$ *for each fixed* $\boldsymbol{\gamma} \in \Gamma$.
2. $\hat{\boldsymbol{\beta}}(\boldsymbol{\gamma}) \xrightarrow{p} \boldsymbol{\beta}_0$ *uniformly over* $\boldsymbol{\gamma} \in \Gamma$.
3. *Under* H_1^*, $\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_0 = O_p(n^{-1})$ *and* $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \xrightarrow{d} \mathcal{N}\{\mathbf{0}, \boldsymbol{V}(\boldsymbol{\gamma}_0)\}$.

- See the full paper for the regularity conditions, the construction of $\boldsymbol{V}(\boldsymbol{\gamma})$, and the proof of Theorem 1.
- Under H_0^*, the asymptotic distribution of $\hat{\boldsymbol{\beta}}$ is **non-standard**.

Testing the no-threshold-effect hypothesis

- Testing H_0^* requires the **wild bootstrap** of Hansen (1996), as γ is unidentified and $\hat{\boldsymbol{\beta}}$ is not asymptotically normal under H_0^*.
- Formulate the no-threshold-effect hypothesis H_0^* as a linear parametric restriction:

$$H_0^* : \boldsymbol{R}^* \boldsymbol{\beta} = \boldsymbol{0} \qquad \text{vs.} \qquad H_1^* : \boldsymbol{R}^* \boldsymbol{\beta} \neq \boldsymbol{0}.$$

  where $\boldsymbol{R}^* = (\boldsymbol{I}_{p+1},\ -\boldsymbol{I}_{p+1})$.
- The Wald test statistic conditional on $\boldsymbol{\gamma}$ is given by:

$$\mathcal{W}_n^*(\boldsymbol{\gamma}) = n\hat{\boldsymbol{\beta}}(\boldsymbol{\gamma})^\top (\boldsymbol{R}^*)^\top \left\{\boldsymbol{R}^* \hat{\boldsymbol{V}}_n(\boldsymbol{\gamma})(\boldsymbol{R}^*)^\top\right\}^{-1} \boldsymbol{R}^* \hat{\boldsymbol{\beta}}(\boldsymbol{\gamma}).$$

- See the full paper for the construction of $\hat{\boldsymbol{V}}_n(\boldsymbol{\gamma})$.

Testing the no-threshold-effect hypothesis

- Incorporate all possible values of γ as in:

$$\text{sup}\mathcal{W}_n^* = \sup_{\boldsymbol{\gamma}\in\Gamma} \mathcal{W}_n^*(\boldsymbol{\gamma}).$$

- Let $g(\mathcal{W}_n^*)$ be either $\text{sup}\mathcal{W}_n^*$, $\text{ave}\mathcal{W}_n^*$, or $\text{exp}\mathcal{W}_n^*$.
- Let $\{g\{\mathcal{W}_n^{*(b)}\}\}_{b=1}^B$ be the set of wild-bootstrap test statistics. (See the full paper for the bootstrap procedure.)
- The bootstrap p-value is defined as:

$$\hat{p}_n^B(H_0^*) = \frac{1}{B}\sum_{b=1}^{B} \mathbf{1}\left[g\left\{\mathcal{W}_n^{*(b)}\right\} \geq g(\mathcal{W}_n^*)\right].$$

- Reject H_0^* if $\hat{p}_n^B(H_0^*) < a$, where $a \in (0,1)$ is the nominal size.

Testing the no-threshold-effect hypothesis

Theorem 2 (Bootstrap test for H_0^*)

Under standard regularity conditions, the following are true:

1. *Under H_0^*, $\hat{p}_n^B(H_0^*)$ is asymptotically uniform on $[0, 1]$.*
2. *Under H_1^*, $\hat{p}_n^B(H_0^*) \xrightarrow{p} 0$ as $n \to \infty$ and $B \to \infty$.*

- See the full paper for the regularity conditions and the proof.
- The bootstrap test for H_0^* is **asymptotically valid**; the test has size approaching the nominal size a under H_0^*, and power approaching 1 under H_1^*.

Midastar II: Motivation

- Now assume the target variable y is observed at a high frequency and the threshold variable x is observed at a low frequency.
- Assume the ratio of sampling frequencies, m, is known and fixed across time. (Example: $m = 3$ if y is sampled monthly and x is sampled quarterly.)
- For each $t \in \mathbb{L}$, we sequentially observe $\{y_j^*\}_{j \in \mathbb{H}_t}$ for the target variable, while we observe a single realization x_t for the threshold variable.
- For $t \in \mathbb{L}$, let $y_t = \sum_{k=1}^{m} w_k y_{t-1+k/m}^*$ be an aggregation of y^*.
- In the existing literature, the single-frequency TAR model is fitted to $\{y_t, x_t\}_{t \in \mathbb{L}}$ even if $\{y_t^*\}_{t \in \mathbb{H}}$ are observable.
- Such a temporal aggregation can make inference less accurate.

Midastar II: Specification

- To avoid the temporal aggregation of y^*, we propose **the Midastar model of the second kind (Midastar II)**:

$$y_t^* = \begin{cases} \alpha_1 + \sum_{k=1}^{p} \phi_{1k} y^*_{t-\frac{k}{m}} + u_t^* & \text{if } \; x_{\lceil t \rceil - d} < \mu, \\ \alpha_2 + \sum_{k=1}^{p} \phi_{2k} y^*_{t-\frac{k}{m}} + u_t^* & \text{if } \; x_{\lceil t \rceil - d} \geq \mu, \quad t \in \mathbb{H}. \end{cases}$$

- Since x is observed at the low frequency, the delay parameter d needs to be integer-valued (i.e., $\mathcal{D} \subseteq \mathbb{N}$).
- Similarly, the choice set of μ needs to be $\mathcal{X} = \{x_t\}_{t \in \mathbb{L}}$.
- Suppose $x_{\lceil t \rceil - d} < \mu$ for some $t \in \mathbb{H}$, then regime 1 keeps arising at the m consecutive high frequency periods:

$$\mathbb{H}_{\lceil t \rceil} = \{\lceil t \rceil - 1 + 1/m, \; \lceil t \rceil - 1 + 2/m, \; \ldots, \; \lceil t \rceil\}.$$

Midastar II: Threshold effects

- The aggregated TAR model with finite lag length $q \in \mathbb{N}$ is generally misspecified relative to the Midastar II, as aggregating an AR(p) process results in an infinite-order AR process in general (e.g., Lütkepohl, 1984; Ghysels, Hill & Motegi, 2016).
- Hence, fitting the aggregated TAR model can lead to two types of incorrect conclusions:
    1. pointing to the absence of threshold effects when present in reality (i.e., spurious non-threshold effects);
    2. pointing to the presence of threshold effects when absent in reality (i.e., spurious threshold effects).
- Spurious non-threshold effects are the major problem of aggregating y^*. It is well known that temporal aggregation is inclined to weaken threshold effects (Granger & Lee, 1999).

Midastar II: Threshold effects

- Spurious threshold effects are a more manageable problem.
- If y^* follows the single-regime AR process, then the aggregated y follows AR(∞) in general.
- Fitting the aggregated TAR model with finite lag length $q \in \mathbb{N}$ results in misspecification, but the original process can be well approximated by using q_n lags such that $q_n \to \infty$ as $n \to \infty$ at a proper rate (e.g., Lütkepohl & Poskitt, 1996).
- Thus, it should be technically feasible to reach the correct conclusion of no threshold effects via the aggregated TAR(q_n).

Midastar II: Key features

- The number of parameters in the Midastar II does not depend on m. Hence, the Midastar II operates well when m takes a large or time-varying value.
- Another virtue of the Midastar models is that both Scenario I (i.e., low frequency y and high frequency x^*) and Scenario II (i.e., high frequency y^* and low frequency x) are easy to handle.
- In the existing MIDAS models, Scenario II is less straightforward than Scenario I (e.g., Ghysels, Sinko & Valkanov, 2007; Ghysels, Hill & Motegi, 2020).

Monte Carlo simulation I: Design

- Suppose the DGP is the Midastar I:

$$y_t = \begin{cases} \alpha_{10} + \phi_{10} y_{t-1} + \epsilon_t & \text{if } x^*_{t-\frac{d_0}{m}} < \mu_0, \\ \alpha_{20} + \phi_{20} y_{t-1} + \epsilon_t & \text{if } x^*_{t-\frac{d_0}{m}} \geq \mu_0, \end{cases} \quad t \in \mathbb{L}.$$

- Suppose $\alpha_{10} = \alpha_{20} = 0$, $d_0 = 1$, $\mu_0 = 0$, and $\epsilon_t \overset{i.i.d.}{\sim} \mathcal{N}(0,1)$.
- Let $\phi_{10} = 0.2$ and $\phi_{20} \in \{0.2, 0.8\}$.
- When $\phi_{20} = 0.2$, threshold effects are absent and the DGP reduces to the single-regime AR(1) process.
- When $\phi_{20} = 0.8$, threshold effects are present and the DGP does not degenerate.

Monte Carlo simulation I: Design

- Suppose the DGP of x^* is AR(1):

$$x_t^* = 0.4 x_{t-\frac{1}{m}}^* + \nu_t^*, \quad \nu_t^* \overset{i.i.d.}{\sim} \mathcal{N}(0,1), \quad t \in \mathbb{H}.$$

- Assume $\{\epsilon_t\}_{t \in \mathbb{L}}$ and $\{\nu_t^*\}_{t \in \mathbb{H}}$ are mutually independent.
- Let $m = 3$ and $n \in \{120, 240, 480\}$.
- The Midastar model of the first kind is specified as

$$y_t = \begin{cases} \alpha_1 + \phi_1 y_{t-1} + u_t & \text{if } \ x_{t-\frac{d}{m}}^* < \mu, \\ \alpha_2 + \phi_2 y_{t-1} + u_t & \text{if } \ x_{t-\frac{d}{m}}^* \geq \mu, \quad t \in \mathbb{L}. \end{cases}$$

- The space of the delay parameter d is $\mathcal{D} = \{1, 2, 3\}$.
- The space of the threshold parameter μ is $\mathcal{X}^* = \{x_t^*\}_{t \in \mathbb{H}}$.

Monte Carlo simulation I: Design

- For each of $J = 1000$ Monte Carlo samples, we estimate $\boldsymbol{\beta} = (\alpha_1, \phi_1, \alpha_2, \phi_2)^\top$ and $\boldsymbol{\gamma} = (d, \mu)^\top$ via profiling.
- The null hypothesis of no threshold effects is written as

$$H_0^* : \phi_{10} = \phi_{20} \qquad \text{vs.} \qquad H_1^* : \phi_{10} \neq \phi_{20}.$$

- To test H_0^*, we perform the bootstrap with the sup-Wald, ave-Wald, and exp-Wald statistics as well as their LM versions.
- The rejection frequencies of each test are computed, where the nominal size is 5% and the number of bootstrap iterations is $B = 500$.
- The rejection frequencies correspond to empirical **size** when $\phi_{10} = \phi_{20}$ and empirical **power** when $\phi_{10} \neq \phi_{20}$.

Monte Carlo simulation I: Design

- To inspect the consequence of temporal aggregation, we also fit the TAR model after aggregating x^*:

$$y_t = \begin{cases} \alpha_1 + \phi_1 y_{t-1} + u_t & \text{if } \ x_{t-d} < \mu, \\ \alpha_2 + \phi_2 y_{t-1} + u_t & \text{if } \ x_{t-d} \geq \mu, \quad t \in \mathbb{L}. \end{cases}$$

- The aggregation scheme is the stock aggregation:

$$x_t = x_t^*, \quad t \in \mathbb{L}.$$

- The space of the delay parameter d is $\mathcal{D} = \{1, 2, 3\}$.
- The space of the threshold parameter μ is $\mathcal{X} = \{x_t\}_{t \in \mathbb{L}}$.

Monte Carlo simulation I: Results

Performance of the profiling estimation (Midastar I, $\phi_{20} = 0.2$)

	$n = 120$		$n = 240$		$n = 480$	
	Bias	Stdev	Bias	Stdev	Bias	Stdev
α_1	0.006	0.269	-0.002	0.190	0.002	0.133
ϕ_1	-0.009	0.274	-0.007	0.170	-0.007	0.126
α_2	0.001	0.260	0.006	0.193	0.000	0.128
ϕ_2	-0.006	0.251	0.006	0.171	-0.004	0.127
d	0.995	0.820	1.014	0.829	0.948	0.808
μ	-0.015	0.750	-0.016	0.730	0.010	0.726

- In this table, $\boldsymbol{\gamma}$ is unidentified since $\phi_{10} = \phi_{20} = 0.2$.
- $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 = O_p(n^{-1/2})$ (**consistency**).
- $\hat{\boldsymbol{\gamma}} = O_p(1)$ (**inconsistency**).
- These results are in line with our theorems.

Monte Carlo simulation I: Results

Performance of the profiling estimation (Midastar I, $\phi_{20} = 0.8$)

	$n = 120$		$n = 240$		$n = 480$	
	Bias	Stdev	Bias	Stdev	Bias	Stdev
α_1	0.001	0.167	−0.002	0.105	0.000	0.066
ϕ_1	−0.032	0.152	−0.019	0.077	−0.009	0.053
α_2	0.011	0.173	0.005	0.100	0.002	0.067
ϕ_2	−0.013	0.156	−0.001	0.079	−0.002	0.053
d	0.134	0.449	0.006	0.089	0.000	0.000
μ	0.004	0.346	−0.001	0.149	−0.001	0.070

- In this table, $\boldsymbol{\gamma}$ is identified since $\phi_{10} \neq \phi_{20}$.
- $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 = O_p(n^{-1/2})$ (**consistency**).
- $\hat{\boldsymbol{\gamma}} - \boldsymbol{\gamma}_0 = O_p(n^{-1})$ (**super-consistency**).
- These results are in line with our theorems.

Monte Carlo simulation I: Results

Rejection frequencies of the bootstrap tests for H_0^* (Midastar I)

	$\phi_{20} = 0.2$ (empirical size)			$\phi_{20} = 0.8$ (empirical power)		
	$n = 120$	$n = 240$	$n = 480$	$n = 120$	$n = 240$	$n = 480$
sup-Wald	0.251	0.151	0.095	0.876	0.982	1.000
ave-Wald	0.092	0.076	0.069	0.764	0.962	1.000
exp-Wald	0.188	0.117	0.079	0.863	0.983	1.000
sup-LM	0.029	0.037	0.046	0.562	0.952	1.000
ave-LM	0.028	0.047	0.055	0.527	0.911	1.000
exp-LM	0.024	0.040	0.040	0.600	0.962	1.000

- For all tests, the empirical size approaches 5% as $n \to \infty$.
- The LM tests have sharper size than the Wald tests in small samples.
- For all tests, the empirical power is sufficiently high in small samples, and approaches 100% as $n \to \infty$.

Monte Carlo simulation I: Results

Performance of the profiling estimation of the aggregated TAR (DGP: Midastar I, $\phi_{20} = 0.2$)

	$n = 120$		$n = 240$		$n = 480$	
	Bias	Stdev	Bias	Stdev	Bias	Stdev
α_1	-0.011	0.262	0.005	0.188	-0.001	0.133
ϕ_1	-0.009	0.251	-0.006	0.177	-0.000	0.126
α_2	0.002	0.263	0.003	0.190	0.002	0.127
ϕ_2	-0.014	0.254	-0.009	0.178	0.001	0.126

- In this table, the true DGP is the single-regime AR(1).
- Hence, the aggregated TAR model is correctly specified.
- $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 = O_p(n^{-1/2})$ (**consistency**).

Monte Carlo simulation I: Results

Performance of the profiling estimation of the aggregated TAR (DGP: Midastar I, $\phi_{20} = 0.8$)

	$n = 120$		$n = 240$		$n = 480$	
	Bias	Stdev	Bias	Stdev	Bias	Stdev
α_1	0.003	0.269	-0.001	0.180	0.006	0.131
ϕ_1	0.252	0.271	0.256	0.183	0.266	0.137
α_2	0.001	0.279	-0.002	0.195	-0.001	0.129
ϕ_2	-0.308	0.247	-0.281	0.184	-0.275	0.129

- In this table, the true DGP is the two-regime Midastar I.
- Hence, the aggregated TAR model is **misspecified**.
- $\hat{\phi}_1$ is positively biased and $\hat{\phi}_2$ is negatively biased.
- Both $\hat{\phi}_1$ and $\hat{\phi}_2$ are converging to around 0.5, a signal of **spurious non-threshold effects**.

Monte Carlo simulation I: Results

Rejection frequencies of the bootstrap LM tests for H_0^*

based on the aggregated TAR (DGP: Midastar I)

	$n = 120$			$n = 240$			$n = 480$		
ϕ_{20}	sup	ave	exp	sup	ave	exp	sup	ave	exp
0.2	0.030	0.040	0.031	0.045	0.048	0.040	0.044	0.039	0.041
0.8	0.030	0.040	0.027	0.034	0.049	0.041	0.067	0.081	0.072

- When $\phi_{20} = 0.2$, the true DGP is the single-regime AR(1).
    1. The aggregated TAR model is correctly specified.
    2. The empirical size is sufficiently close to 5%.
- When $\phi_{20} = 0.8$, the true DGP is the two-regime Midastar I.
    1. The aggregated TAR model is misspecified.
    2. There is **almost no power** (spurious non-threshold effects).

Empirical application I: Set-up

- There is a rapidly growing literature in which time series analysis is performed on COVID-19 statistics (e.g., Motegi, Dennis & Hamori, 2022).
- We analyze the threshold effects of the number of new confirmed cases ($Case$) on the number of patients in hospital ($Hosp$) in Japan.
- Our World in Data (OWID) is a well-known data source for the COVID-19 research.
- For Japan, OWID offers **daily** data of $Case$ and **weekly** data of $Hosp$. This motivates the use of the Midastar I with $m = 7$, where the target variable is weekly $Hosp$ and the threshold variable is daily $Case$.

Empirical application I: Set-up

- Here are some notations:

$$
\begin{aligned}
Hosp_t &= \text{the number of COVID-19 patients in hospital} \\
&\quad\ \text{per million people at week } t \in \mathbb{L}; \\
Case_t &= \text{the number of new confirmed cases of COVID-19} \\
&\quad\ \text{per million people at day } t \in \mathbb{H}; \\
y_t &= \Delta \ln Hosp_t \\
&= \ln Hosp_t - \ln Hosp_{t-1} \\
&= \text{the weekly growth of } Hosp; \\
x_t^* &= \Delta \ln Case_t \\
&= \ln Case_t - \ln Case_{t-1/m} \\
&= \text{the daily growth of } Case.
\end{aligned}
$$

Empirical application I: Data

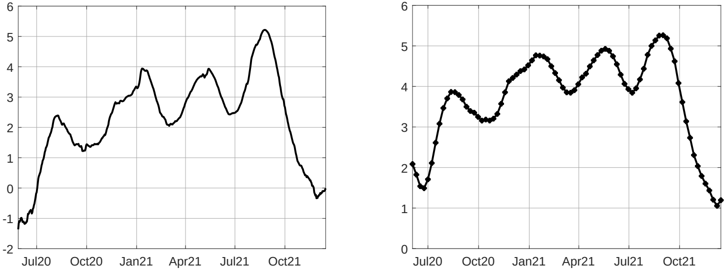

$\ln Case$ (day) $\ln Hosp$ (week)

- Sample period: May 28, 2020 – December 15, 2021.
- Sample size: $n = 81$ weeks, or $mn = 567$ days.
- We observe the second through fifth waves of the pandemic.

Empirical application I: Data

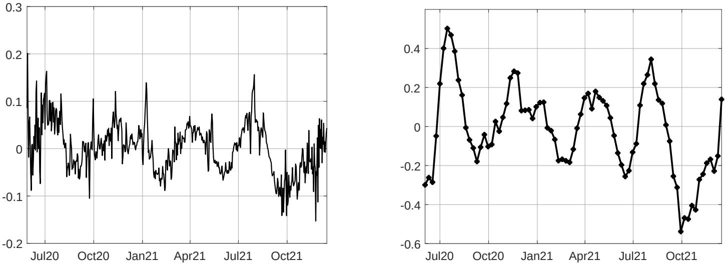

$x^* = \Delta \ln Case$ (day) $\qquad y = \Delta \ln Hosp$ (week)

- Are there threshold effects of $\Delta \ln Case$ on $\Delta \ln Hosp$?
- Does the growth of the number of the patients in hospital have heterogeneous autocorrelation structures when the number of new confirmed cases is below versus above a certain threshold?

Empirical application I: Methodology

- We fit the Midastar I with $p = 4$ and $m = 7$:

$$y_t = \begin{cases} \alpha_1 + \sum_{k=1}^{4} \phi_{1k} y_{t-k} + u_t & \text{if } \; x^*_{t-d/7} < \mu, \\ \alpha_2 + \sum_{k=1}^{4} \phi_{2k} y_{t-k} + u_t & \text{if } \; x^*_{t-d/7} \geq \mu, \end{cases} \quad t \in \mathbb{L}.$$

- Regime 1: contraction phase. Regime 2: expansion phase.
- The space of the delay parameter d is $\mathcal{D} = \{1, \ldots, 21\}$.
- The space of the threshold parameter μ is $\mathcal{X}^* = \{x^*_t\}_{t \in \mathbb{H}}$.
- Let $\boldsymbol{\beta}_r = (\alpha_r, \phi_{r1}, \ldots, \phi_{r4})^\top$ for regime $r \in \{1, 2\}$.
- Estimate $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \boldsymbol{\beta}_2^\top)^\top$ and $\boldsymbol{\gamma} = (d, \mu)^\top$ via profiling.
- Test the no-threshold-effect hypothesis $H_0^* : \boldsymbol{\beta}_1 = \boldsymbol{\beta}_2$ based on the bootstrap exp-LM test.

Empirical application I: Results

- The estimated threshold parameter is $\hat{\mu} = 0.001$.
- The share of the contraction phase to the whole sample is exactly 50%.
- The estimated delay parameter is $\hat{d} = 8$ days.
- The no-threshold-effect hypothesis H_0^* is **rejected** at the 5% level, with the bootstrap p-value being 0.022. Hence, the persistence structures of the number of patients in hospital differ significantly across the regimes.

Empirical application I: Results

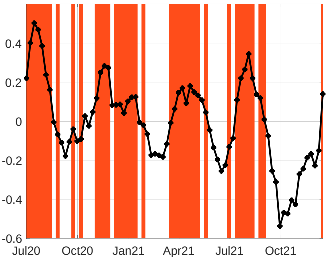

y (shade: regime 2)

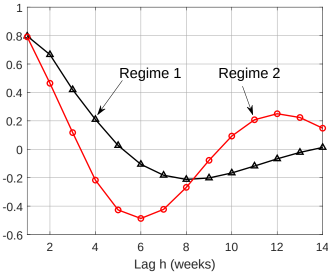

Implied autocorrelations

- We roughly observe that $y_t < 0$ when $x^*_{t-\hat{d}/7} < 0$ (i.e., contraction) and $y_t > 0$ when $x^*_{t-\hat{d}/7} \geq 0$ (i.e., expansion).
- The regime-specific sample mean of y is -0.049 for the contraction phase and 0.028 for the expansion phase.

Empirical application I: Results

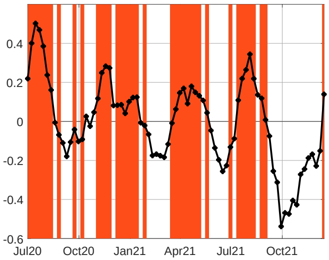

y (shade: regime 2)

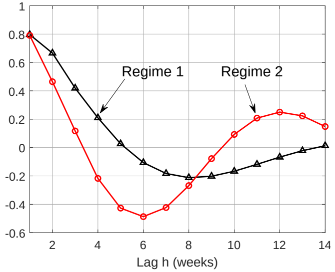

Implied autocorrelations

- The regime-wise implied autocorrelations $\hat{\rho}_r(h)$ are plotted.
- The oscillation of the correlation under the expansion has the larger amplitude and the higher frequency than under the contraction.

Empirical application I: Results

- For comparison, we fit the aggregated TAR model with $p = 4$.
- The estimated delay parameter is $\hat{d} = 1$ week, almost agreeing with our main result of $\hat{d} = 8$ days.
- The no-threshold-effect hypothesis H_0^* **cannot be rejected** at the 5% level, with the bootstrap p-value of the exp-LM test being 0.076.
- The Midastar I outperforms the aggregated TAR model in terms of detecting the threshold effects.

Empirical application II: Summary

- In the full paper, we fit the Midastar II to U.S. macroeconomic indicators:
    1. $y^* =$ monthly nonfarm employment.
    2. $x =$ quarterly real gross domestic product.
- Sample period: 1972M1–2007M12 ($n = 144$ quarters, or $mn = 432$ months).
- The bootstrap exp-LM test based on the Midastar II **rejects** H_0^*, with the p-value being 0.029. Hence, the persistence structures of the employment differ significantly across the regimes.
- The bootstrap exp-LM test based on the aggregated TAR **fails to reject** H_0^*, with the p-value being 0.371. This result suggests averaging y^* has weaken the threshold effects.

Conclusion

- We have proposed the **Midastar** models, a novel extension of TAR to mixed frequency data.
- The Midastar I is designed for the low frequency target variable y and the high frequency threshold variable x^*.
- The Midastar II is designed for the high frequency target variable y^* and the low frequency threshold variable x.
- The Midastar models are capable of detecting **threshold effects** accurately, since all observations available are fully exploited.
- The aggregated TAR can point to **spurious non-threshold effects** due to the loss of information.

Conclusion

- The regression and nuisance parameters of the Midastar models can be estimated via profiling.
- The no-threshold-effect hypothesis can be tested by the wild-bootstrap tests of Hansen (1996).
- The proposed estimation and testing satisfy desired statistical properties in both large and small samples.
- The Midastar I is applied to the COVID-19 data of Japan, finding **significant** threshold effects.
- The Midastar II is applied to the U.S. macroeconomic indicators, finding **significant** threshold effects.
- These significant threshold effects **vanish** once the temporal aggregation is executed (spurious non-threshold effects).

References

- Foroni, C., M. Marcellino, and C. Schumacher (2015). Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials. *Journal of Royal Statistical Society, Series A*, vol. 178, pp. 57–82.
- Ghysels, E. (2016). Macroeconomics and the reality of mixed frequency data. *Journal of Econometrics*, vol. 193, pp. 294–314.
- Ghysels, E., J. B. Hill, and K. Motegi (2016). Testing for Granger causality with mixed frequency data. *Journal of Econometrics*, vol. 192, pp. 207–230.
- Ghysels, E., J. B. Hill, and K. Motegi (2020). Testing a large set of zero restrictions in regression models, with an application to mixed frequency Granger causality. *Journal of Econometrics*, vol. 218, pp. 633–654.

References

- Ghysels, E., P. Santa-Clara, and R. Valkanov (2004). The MIDAS Touch: Mixed Data Sampling Regression Models. Working paper at UCLA.
- Ghysels, E., A. Sinko, and R. Valkanov (2006). Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, vol. 131, pp. 59–95.
- Granger, C. W. J. and T.-H. Lee (1999). The effect of aggregation on nonlinearity. *Econometric Reviews*, vol. 18, pp. 259–269.
- Granger, C. W. J. and T. Teräsvirta (1993). *Modelling Nonlinear Economic Relationships*. Oxford University Press.
- Hansen, B. E. (1996). Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica*, vol. 64, pp. 413–430.

References

- Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, vol. 68, pp. 575–603.
- Hansen, B. E. (2017). Regression kink with an unknown threshold. *Journal of Business & Economic Statistics*, vol. 35, pp. 228–240.
- Lütkepohl, H. (1984). Linear transformations of vector ARMA processes. *Journal of Econometrics*, vol. 26, pp. 283–293.
- Lütkepohl, H. and D. S. Poskitt (1996). Testing for causation using infinite order vector autoregressive processes. *Econometric Theory*, vol. 12, pp. 61–87.

References

- Motegi, K., X. Cai, S. Hamori, and H. Xu (2020). Moving average threshold heterogeneous autoregressive (MAT-HAR) models. *Journal of Forecasting*, vol. 39, pp. 1035–1042.
- Motegi, K., J. W. Dennis, and S. Hamori (2022). Conditional Threshold Autoregression (CoTAR). Working paper at Kobe University and IDA.
- Silvestrini, A. and D. Veredas (2008). Temporal aggregation of univariate and multivariate time series models: A survey. *Journal of Economic Surveys*, vol. 22, pp. 458–497.
- Tong, H. (1978). On a threshold model. In: *Pattern Recognition and Signal Processing*, ed. by C. H. Chen. Sijthoff & Noordhoff, Amsterdam.