BANK OF CANADA
BANQUE DU CANADA

Staff analytical paper / Document analytique du personnel—2026-43

Last updated: September 18, 2026

Segmented Inflation Dynamics (SID): Identifying Trend Inflation from the Price Index

André Binette
Canadian Economic Analysis Department
Bank of Canada
abinette@bank-banque-canada.ca

Colleen Smith
Canadian Economic Analysis Department
Bank of Canada
csmith@bank-banque-canada.ca

DOI: https://doi.org/10.34989/sap-2026-43 | ISSN 3110-9195

Acknowledgements

We thank colleagues at the Bank of Canada for valuable comments and suggestions. We are especially grateful to Alexander Ueberfeldt for helpful discussions and detailed feedback as this work developed. Any remaining errors and omissions are the author's sole responsibility.

Abstract

This paper presents Segmented Inflation Dynamics (SID), a data-driven tool for assessing trend inflation. SID uses time-series segmentation to divide the price index into contiguous segments with stable inflation rates. The segmentation breaks are identified from the data. The slope of each segment provides a well-defined estimate of trend inflation that is stable but responsive to persistent shifts. Using Canadian data from 1992 to 2025, the paper describes the resulting inflation segments, tests the robustness of the method and assesses its real-time properties. The results show that SID offers an interpretable measure of trend inflation, with limited sensitivity to short-term volatility and a quick adjustment to shifts in inflation dynamics. Based on data through December 2025, SID indicates that trend inflation in Canada has been close to 2 percent since early 2024.

Topics: Inflation; Trend inflation; Core inflation; Time-series segmentation; Monetary policy
JEL codes: E31; C22; E52

Résumé

Cette étude présente la dynamique segmentée de l'inflation, ou SID, un outil fondé sur les données pour évaluer l'inflation tendancielle. SID utilise la segmentation de séries chronologiques pour diviser l'indice des prix en segments contigus caractérisés par des taux d'inflation stables. Les points de rupture entre les segments sont déterminés à partir des données. La pente de chaque segment fournit une estimation bien définie de l'inflation tendancielle, à la fois stable et sensible aux changements persistants. À partir de données canadiennes couvrant la période de 1992 à 2025, cette étude décrit les segments d'inflation obtenus, teste la robustesse de la méthode et évalue ses propriétés en temps réel. Les résultats montrent que SID fournit une mesure interprétable de l'inflation tendancielle, peu sensible à la volatilité de court terme et capable de s'ajuster rapidement aux changements dans la dynamique de l'inflation. Selon les données disponibles jusqu'en décembre 2025, SID indique que l'inflation tendancielle au Canada se situe près de 2 % depuis le début de 2024.

Sujets : Inflation; Inflation tendancielle; Inflation fondamentale; Segmentation de séries chronologiques; Politique monétaire
Codes JEL : E31; C22; E52

1. Introduction

Distinguishing persistent inflation pressures from temporary price movements is essential for monetary policy. Measures of trend inflation play a critical role in this task, yet estimating trend inflation in real time remains challenging. Trend inflation helps interpret movements in headline inflation by looking through temporary fluctuations caused by volatile commodity prices, tax changes, and one-off supply shocks. The challenge is even greater when large relative-price changes, supply disruptions or shifts in demand patterns affect inflation.

No single measure perfectly captures trend inflation, so central banks rely on a range of complementary indicators. These include exclusion-based measures, trimmed-mean and median estimators, persistence-weighted indicators, and common-trend or factor-model approaches (Gordon 1975; Bryan and Cecchetti 1994; Cutler 2001; Reis and Watson 2010; Stock and Watson 2016; Uzeda 2025).[1] These measures are interpreted alongside information on the distribution of price change, sectoral shocks and forward-looking information—most notably inflation expectations—to form a multifaceted view of underlying inflation (Mendez 2025).

Most existing approaches estimate trend inflation from transformations of the price index, such as monthly or yearly inflation rates, sometimes incorporating additional data or using complex techniques. This paper introduces Segmented Inflation Dynamics (SID), a complementary approach that extracts trend inflation directly from the price index. SID partitions the series into contiguous segments with stable inflation rates. The slope of each segment represents the estimated trend inflation during that period, and once annualized, is directly comparable to conventional indicators. This technique, known as time-series segmentation, originates in data science and, to our knowledge, has not previously been applied to measuring trend inflation.

SID allows the data to determine the relevant time scale, rather than relying on fixed windows. More fundamentally, SID views trend inflation as a sequence of stable periods in the evolution of the price index. While related to the literature on regime-switching and structural break models, SID does not impose a predetermined number of states or parametric transition dynamics. Instead, break points emerge directly from the data. As a results, SID provides a stable representation of inflation pressures while remaining responsive to shifts in inflation dynamics.

Using Canadian data, we show how SID works and document the implied trend inflation rates since 1992. The results indicate that SID produces an interpretable estimate of

[1] In the early years of inflation targeting, the Bank of Canada relied on the consumer price index (CPI) excluding food and energy (Laflèche, 1997). In 2001, the Bank adopted CPIX, which excludes eight volatile components and the effect of indirect taxes (Macklem, 2001). In 2016, the Bank introduced CPI-trim, CPI-median and CPI-common as its preferred core measures (Khan & al., 2015). The Bank later removed CPI-common from this set because of large and persistent historical revisions reduced its usefulness for real-time policy assessment (Sullivan, 2022).

trend inflation, with clear transitions between periods and limited sensitivity to short-term volatility. Based on data through December 2025, SID suggests that trend inflation has been near 2 percent since early 2024. This stability is consistent with inflation having re-anchored around the Bank of Canada's 2 percent target following recent shocks.

The rest of this paper is organized as follows: Section 2 describes the methodology and data. Section 3 discusses results. Section 4 concludes.

2. Methodology and data

Many indicators of trend inflation exist, each with strengths and limitations. Time-series segmentation offers a simple, data-driven alternative designed to capture persistent dynamics in the data. It can be applied directly to the consumer price index (CPI), without exclusions or fixed time intervals. Because inflation often moves through relatively stable periods interrupted by shifts in underlying price pressures, segmenting the price index provides a natural way to track changes in trend inflation. In this paper, we refer to these stable periods as segments to distinguish them from regimes in parametric regime-switching frameworks.

2.1 Segmentation algorithm

Time-series segmentation partitions a series into contiguous, homogeneous segments. The technique is widely used for trend extraction and data compression (Keogh et al. 2004; Lovrić et al. 2014) and has been applied in financial and macroeconomic analysis (Zhang et al. 2007; Cheong et al. 2012; Sousa et al. 2020).

Our segmentation procedure follows the feasibility space (FS) algorithm of Liu, Lin, and Wang (2008). The series is divided into non-overlapping linear[2] segments subject to two constraints:

1. Each segment forms a straight line connecting two observed points.
2. All observations within the segment lie within a tolerance band defined by δ.

With each new observation, the feasible set of slopes is updated. As more points are added, the feasible space narrows, tightening the range of slopes consistent with the segment. When the feasible space becomes empty, a break is identified and a new segment begins. Appendix A illustrates the procedure, and Appendix B provides pseudocode.

The *tolerance parameter* δ controls the sensitivity of the algorithm. A low value of δ generates many short segments that may reflect noise rather than meaningful changes

[2] Linearity within segments is a modelling choice. It reflects the assumption that, over any given period, movements in the price level can be summarized by a constant average rate of change, represented by the segment slope.

in the trend. A high value of δ produces longer segments that may smooth through the data and miss genuine breaks. Choosing δ therefore requires balancing the need to ignore short-term volatility against maintaining sensitivity to sustained shifts in inflation dynamics.

2.2 Dealing with temporary disturbances

The original FS algorithm treats any large deviation (i.e., an outlier) as a potential break. For inflation data, this can lead to spurious breaks caused by volatility in individual CPI components, such as gasoline prices, hotel rates, or communication services. To address this issue, we modify the algorithm so that a break occurs only after a specified number of consecutive outliers. As a results, isolated outliers are ignored.

This additional parameter, together with the tolerance parameter (δ), enables the algorithm to abstract from short-lived disturbances while remaining sensitive to sustained shifts in the data.

2.3 Parameter selection

Candidate specifications are generated through a grid search over two parameters: (i) the tolerance parameter δ and (ii) the maximum number of consecutive outliers, which range from 1 to 4.[3,4]

Selection of the preferred specification follows two steps. First, specifications are screened based on the segmentation structure. A specification is retained only if the share of short segments (S) does not exceed the share of long segments (L). Short segments last fewer than four months, while long segments last more than twenty-four

[3] This procedure assumes that the data-generating process remains sufficiently stable over the inflation-targeting period for a single segmentation rule to apply throughout the sample rather than varying across subperiods.

[4] The tolerance grid is defined in log units. The parameter begins at 0.001 (approximately 0.1%) and increases in increments of 0.0001 (approximately 0.01%) to a maximum of 0.01 (approximately 1%). This range allows us to assess the sensitivity of the algorithm across a broad set of parameter values.

months.[5,6] This restriction avoids specifications dominated by short-lived segments that are more likely to reflect noise rather than meaningful shifts in inflation dynamics.

Second, among the retained candidates, the preferred specification minimizes the root mean squared error (RMSE) between the SID estimate and the annualized month-over-month inflation rate.[7] Using month-over-month inflation as a reference provides a high-frequency benchmark for segmentation fit.

This two-step procedure separates the economic features of the segmentation from the statistical assessment of fit. It ensures the selected specification produces economically meaningful segmentation while maintaining a strong representation of inflation dynamics.

However, there are limitations. First, segmentation is well suited to identifying discrete shifts but may be less effective when trend inflation changes gradually rather than through distinct breaks. Second, results depend on parameter choices, although this sensitivity is mitigated through parameter selection and robustness checks. Third, like other high-frequency approaches, the method can be affected by data noise, particularly at the beginning of new segments.

We assess robustness in Section 3 by comparing segmentation structures and implied trend inflation paths across the retained specifications.

2.4 Measuring trend inflation

Applying the segmentation procedure with the selected value of δ and the consecutive-outliers thresholds yields a piecewise linear representation of the price index.[8] Segment slopes are expressed in annualized terms, allowing direct comparison with headline inflation and conventional measures of trend inflation. Unlike year-over-year inflation,

---

[5] Monthly observations are noisy. Under standard signal-extraction theory, the standard error of a k-month average declines with $1/\sqrt{k}$. As a result, segments lasting at least four months reduce noise by roughly half relative to a single observation and provide a more informative signal of a change in trend. The 24-month cutoff for a long segment is intended to identify segmentations that are overly coarse for inflation analysis. Because inflation is typically assessed over 12-month horizons, segments lasting more than 24 months imply that the inflation rate is broadly unchanged over two full annual cycles. A large share of such segments may indicate that the specification smooths the data excessively and fails to identify meaningful shifts in the inflation environment.

[6] The results are sensitive to these thresholds. For example, defining short segments as those lasting fewer than six months and long segments as those lasting more than eighteen months yields somewhat different outcomes, while other combinations tend to produce more pronounced differences.

[7] We use the RMSE because it is a scale-sensitive measure of fit that penalizes large deviations. The results are robust to using mean absolute error (MAE) instead.

[8] All calculations are performed using the logarithm of the price index, so linear slopes correspond to constant inflation rates.

the SID measure is not tied to a fixed window because segment lengths are determined endogenously.

As with most measures of trend inflation, estimates near the end of the sample may be revised as new observations become available. However, revisions are confined to the most recent unfinished segment; slopes for completed segments remain unchanged. Section 3 examine the method's real-time properties by tracking how segment estimates evolve as new data are added.

2.5 Data: CPI excluding the effect of indirect taxes

The analysis uses monthly, seasonally adjusted CPI data excluding the effect of indirect taxes (**Chart 1**).

Chart 1: CPI inflation and selected core inflation measures

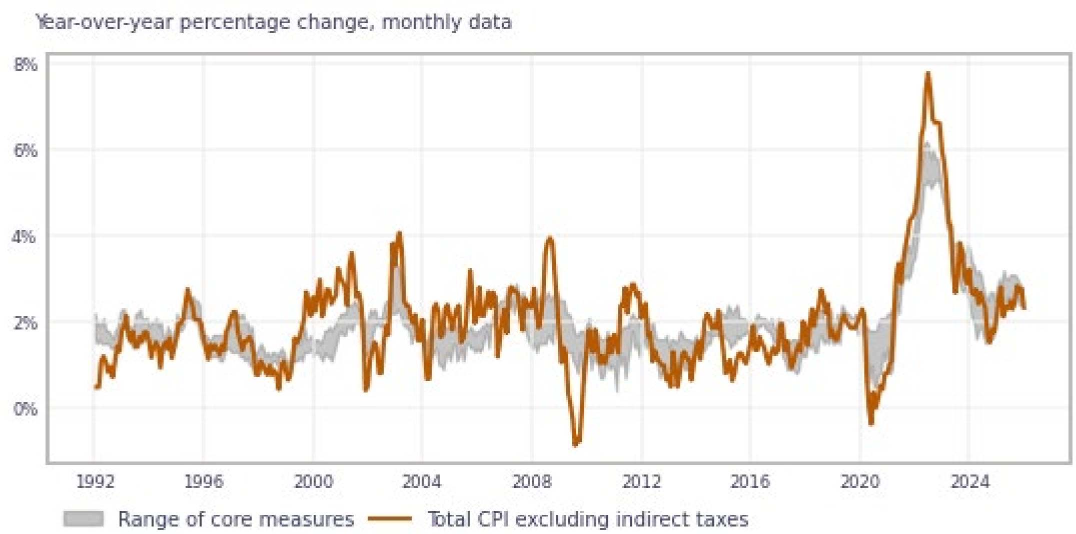

Note: Core measures included are CPI-trim, CPI-median, CPIX and CPIXFET. CPI-trim excludes CPI components whose rates of change in a given month are located in the tails of the distribution of price changes. CPI-median is a measure of core inflation corresponding to the price change located at the 50th percentile. CPIX excludes eight of the most volatile components of inflation. CPIXFET excludes food and energy inflation. All core inflation measures exclude the effect of indirect taxes.
Sources: Statistics Canada and Bank of Canada calculations
Last observation: December 2025

Working directly with the price index improves interpretability, while seasonal adjustment removes predictable fluctuations. Excluding indirect taxes is important because the segmentation procedure is highly sensitive to sudden, large changes in the price level. Changes in sales or excise taxes can create abrupt shifts that the algorithm may interpret as new segments, resulting in short-lived segments with little relevance for trend inflation (see Appendix C).[9]

[9] Major indirect tax changes during the sample include tobacco tax adjustments, GST reductions, sales-tax harmonization, provincial tax measures, the federal GST holiday, and the removal of the federal carbon price on fuels.

3. Results

This section presents results along four dimensions: segmentation outcomes and the implied path of trend inflation, robustness across specifications; real-time behavior and the current assessment of trend inflation.

3.1 Segmentation outcomes and implied trend inflation

Applying SID to CPI excluding the effect of indirect taxes identifies 29 distinct segments between January 1992 and December 2025 (**Figure 1**).[10] Segments last about 12 months on average, although their durations range from one month to several years during periods of stable inflation.

Figure 1: SID – Segmented inflation dynamics

CPI excluding the effect of indirect taxes, log level, monthly data

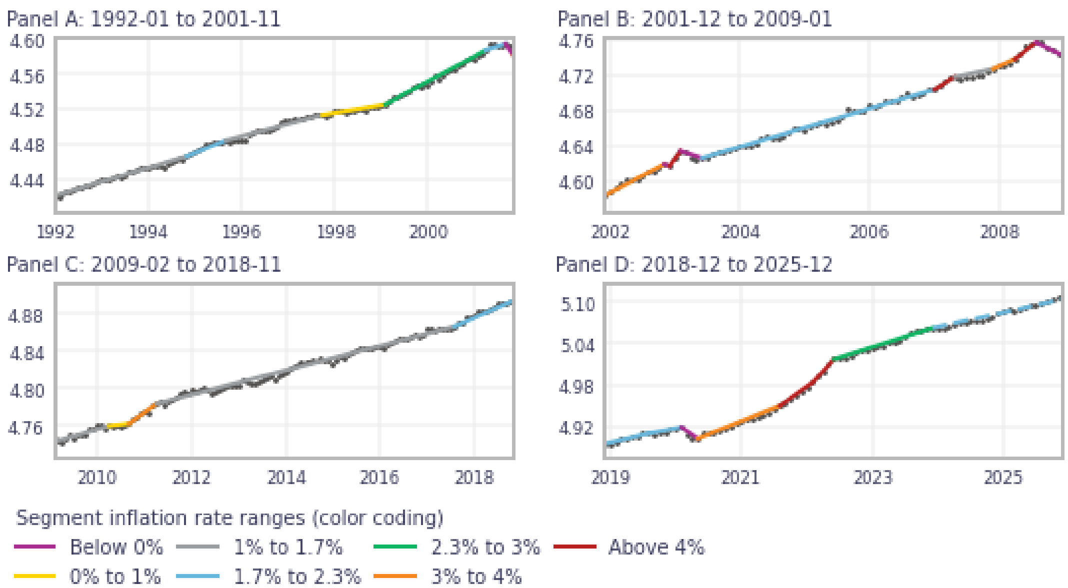

Note: Each segment has a specific estimated inflation rate. For visualization, segments are color-coded into broad inflation ranges.
Sources: Statistics Canada and Bank of Canada calculations
Last observation: December 2025

Short-lived segments often occur around large shocks, including the September 11 attacks, energy price volatility in the early 2000s, the global financial crisis and the onset of the COVID-19 pandemic. For example, SID identifies a one-month segment in December 2002, followed by two additional short segments. These segments reflect an exceptionally volatile period driven by swings in gasoline, natural gas, and electricity prices.

[10] According to the parameter's selection procedure, the preferred specification sets the tolerance parameter δ equal to 0.5 and allows for one outlier before a break is declared.

In contrast, some periods exhibit remarkably stable inflation dynamics. SID identifies a 76-month segment between 2011 and 2017 and a 43-month segment between 2003 and 2007, highlighting prolonged periods of stability. More recently, SID identifies two consecutive high-inflation segments from mid-2021 to mid-2022, with trend inflation near 7 and then 10 percent, before transitioning to lower-inflation segments. The most recent segment, which began in late 2023, indicates that inflation has been trending close to the 2 percent target.

Comparing SID slopes with annualized month-over-month inflation reveals three features (**Figure 2**).

1. Within each segment, the SID slope tracks the central tendency of monthly inflation.
2. Segment changes are infrequent and generally coincide with large, persistent shifts in inflation dynamics.[11]
3. By determining segment lengths endogenously, SID smooths high-frequency volatility and avoids distortions associated with fixed-window growth rates.

Figure 2: Monthly inflation rates versus SID

CPI excluding the effect of indirect taxes, month-over-month percentage change and SID slope, annualized, monthly data

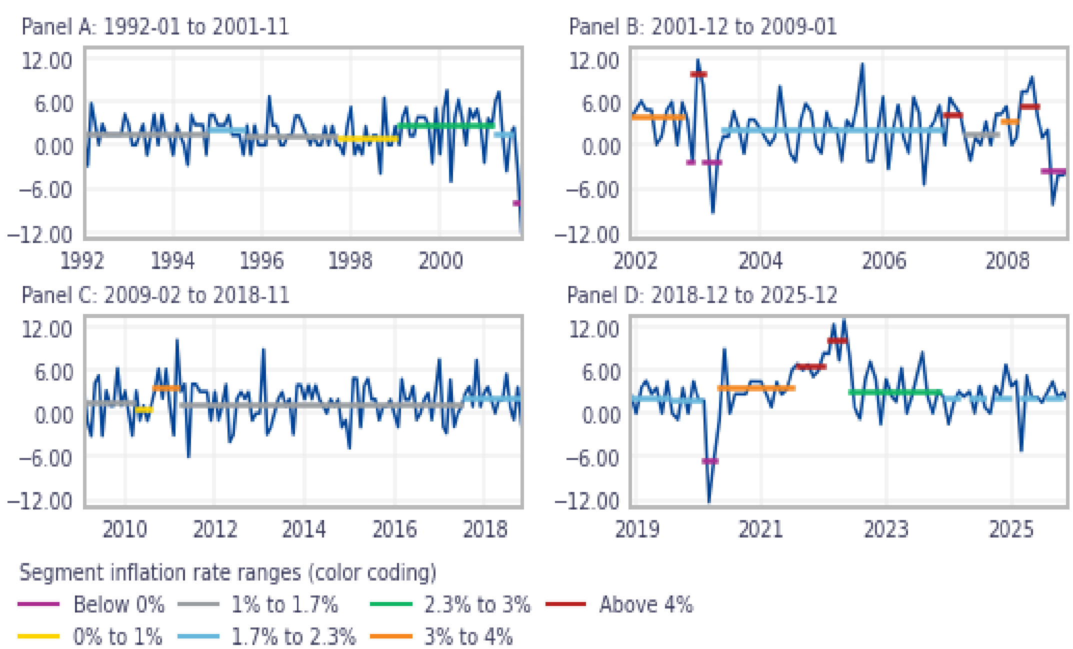

Note: Each segment has a specific estimated inflation rate. For visualization, segments are color-coded into broad inflation ranges.

Sources: Statistics Canada and Bank of Canada calculations

Last observation: December 2025

[11] Although the preferred value of δ implies a relatively wide tolerance band, SID can still detect subtle changes in trend inflation because regime changes are triggered by sustained patterns in the data rather than by isolated large monthly movements.

SID and year-over-year inflation move closely together, but important differences emerge around turning points (**Figure 3**). Because year-over-year inflation is affected by base effects, it can remain elevated or depressed after inflation momentum has shifted. By contrast, SID adjusts as persistent changes in inflation rates emerge. This difference is particularly evident following the pandemic shock. SID identifies a shift to high inflation, followed by a return to a lower trend.

Figure 3: Year-over-year inflation rates versus SID

CPI excluding the effect of indirect taxes, year-over-year percentage change and SID slope annualized, monthly data

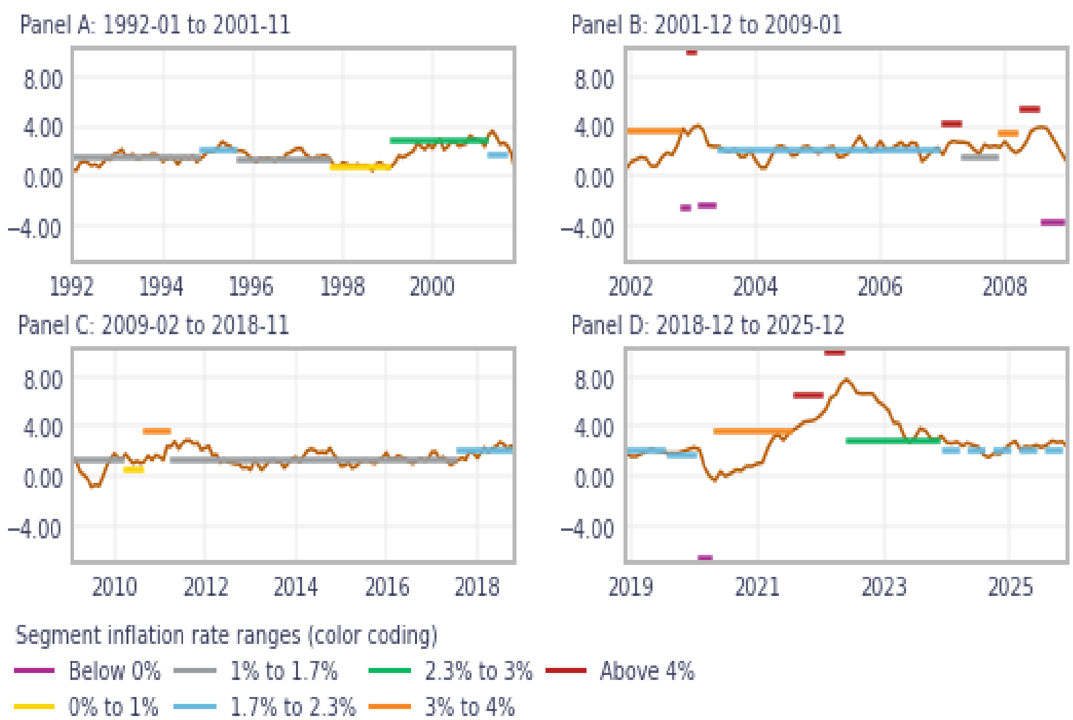

Note: Each segment has a specific estimated inflation rate. For visualization, segments are color-coded into broad inflation ranges.
Sources: Statistics Canada and Bank of Canada calculations
Last observation: December 2025

3.2 Robustness across specifications

We assess robustness across the subset of specifications identified above. The 10 best-performing specifications, ranked by RMSE, produce similar results (**Chart 2**). Segment boundaries for longer-run segments occur at similar dates, and differences are generally confined to periods of elevated volatility.

For two of the top specifications, the slope of the final segment is higher than the others. This reflects the ongoing assessment of post-mid-2023 inflation dynamics, which will evolve as new data become available. As shown in **Chart 3**, once the segment break occurs, trend inflation in these specifications is likely to decline to just above 2 percent, aligning them with the other retained specifications.

Chart 2: SID delivers similar trend inflation signals across top specifications

Segmentations for the ten highest-ranked specifications

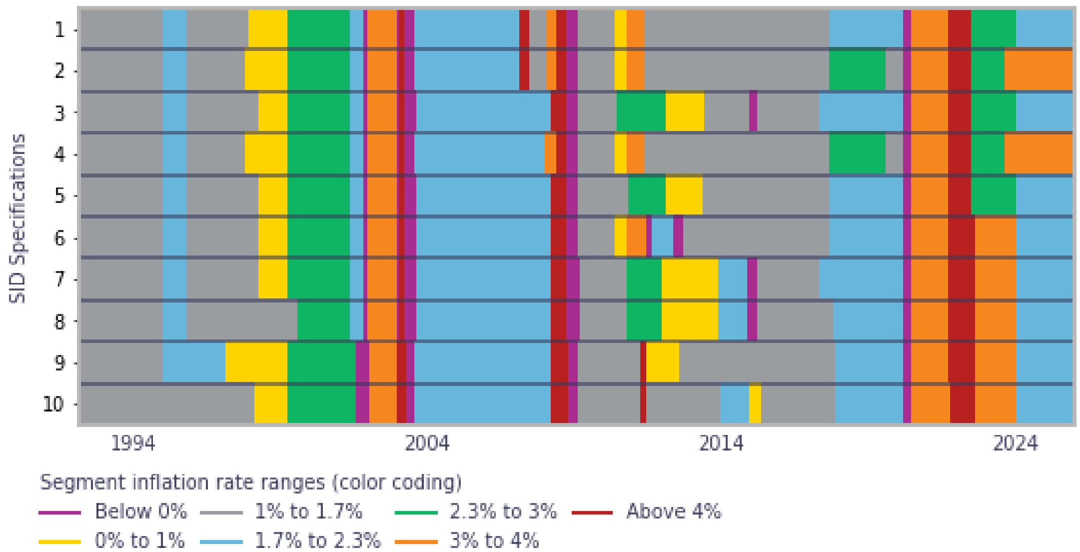

Chart 3: Unfinished segments can mask current trend inflation estimates

CPI excluding the effect of indirect taxes, log level, monthly data

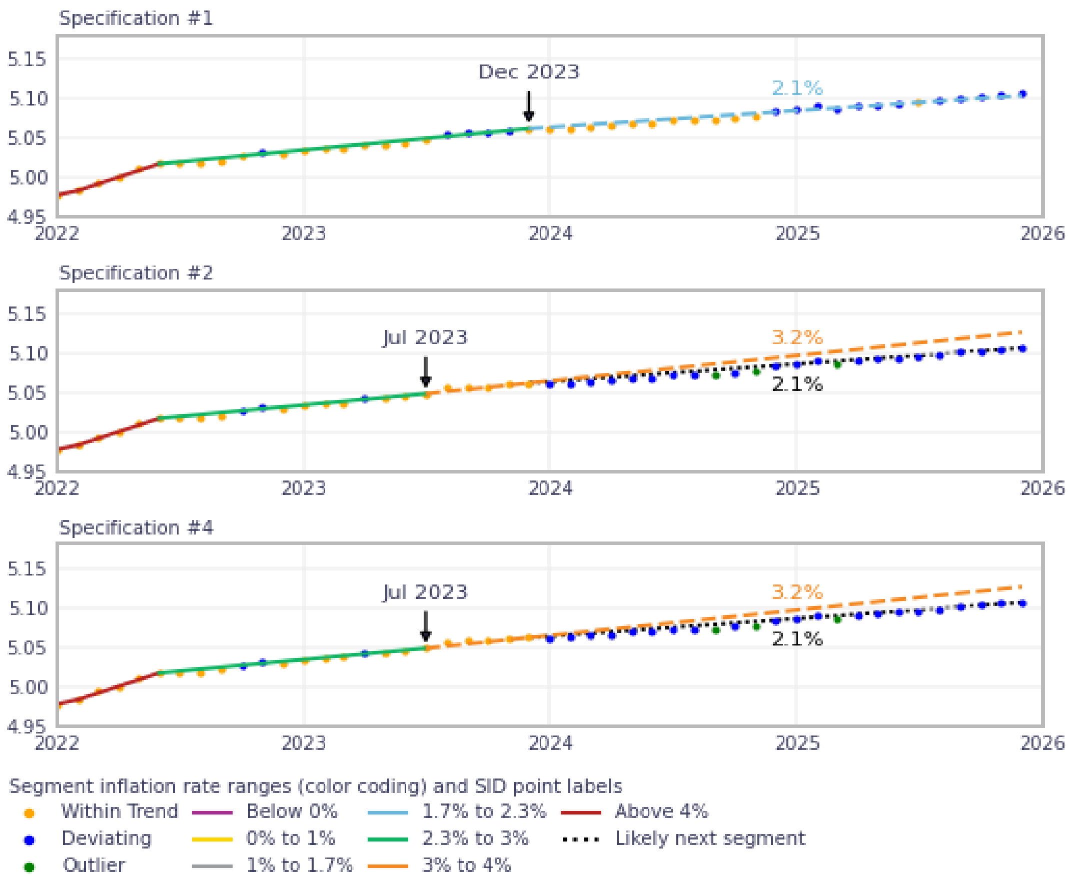

Sources: Statistics Canada and Bank of Canada calculations
Last observation: December 2025

To quantify similarity, we construct an overlap metric that compares each candidate segmentation with the preferred specification. Let $x_A(t)$ denote the SID slope at month t under the preferred specification, and $x_B(t)$ the corresponding slope under an alternative specification. Monthly similarity is defined as:

$$s(t) = exp\left(\frac{-|\ x_A(t) - x_B(t)\ |}{\gamma}\right)$$

where Y defines the scale of meaningful deviations and is equal to the median absolute deviation. This robust measure of dispersion captures typical variation in trend estimates while limiting the influence of extreme values. Differences smaller than Y result in high similarity, differences near Y reduce similarity substantially, and larger differences lead to a rapid decline in similarity.[12]

The overall overlap score is defined as:

$$Overlap = \frac{1}{T}\sum_{t=1}^{T} s(t)$$

where T is the total number of months in the sample. To avoid end-of-sample distortions, we exclude the unfinished segment from the calculation.

The overlap score ranges from 0 to 1. A value of 1 indicates identical segment timing and trend inflation estimates. Lower values reflect greater divergence in timing, levels, or both. We interpret the scores as follows:

- ≥0.85: same qualitative assessment
- 0.70–0.85: limited sensitivity
- 0.50–0.70: material differences
- <0.50: unstable specification

These thresholds are intended as practical benchmarks to facilitate interpretation rather than formal statistical cutoffs.

Most of the 10 top-ranked specifications have overlap scores in the upper range, confirming that they produce similar segment structures and trend inflation estimates (**Chart 4**).[13] This result suggest that the main features of the segmentation—both in timing and in implied trend inflation—are robust to reasonable variations in the parameter values.

[12] The exponential form provides a smooth, bounded similarity score that declines rapidly as differences in slopes increase. This functional form penalizes larger deviations more strongly than linear alternatives while keeping the metric scaled between 0 and 1.

[13] Outside periods of large changes in sales or excise taxes and rebates, the approach yields similar results when applied to total CPI. The overlap score of 0.75 exceeds the threshold for small differences.

Chart 4: High similarity across top specifications

Overlap score, %

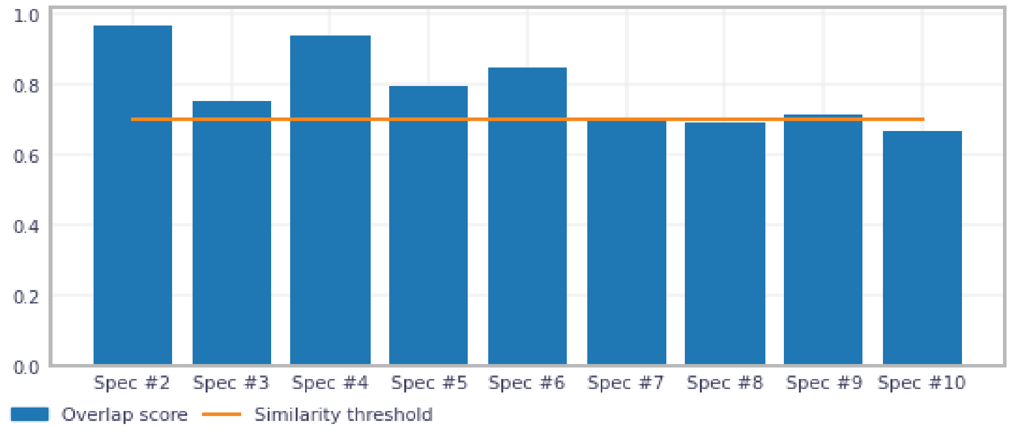

3.3 Real-time behavior

SID produces trend estimates in real time as new CPI observations become available. However, segment endpoints can be confirmed only with additional data and are typically identified about six months after a change in trend. In practice, policymakers may not need to wait that long. Deviations of incoming observations from the prevailing slope often provide early signals of whether trend inflation is picking up or slowing.

Early estimates within a new segment are naturally volatile. During the first three months, the mean absolute deviation from the final trend estimate is about 1 percentage point—or higher—for many segments, particularly those during periods of rapid inflation change (**Chart 5**). Volatility declines quickly thereafter. By month four, deviations fall below 0.5 percentage points on average, and by month six they drop below 0.25 percentage points.

This pattern underscores a key trade-off. SID is designed to respond quickly to incoming data at the start of a segment, which results in higher short-run volatility. As observations accumulate, however, the estimate converges rapidly toward its final value.

Chart 5: Volatility of SID trends falls after 4 and 6 months

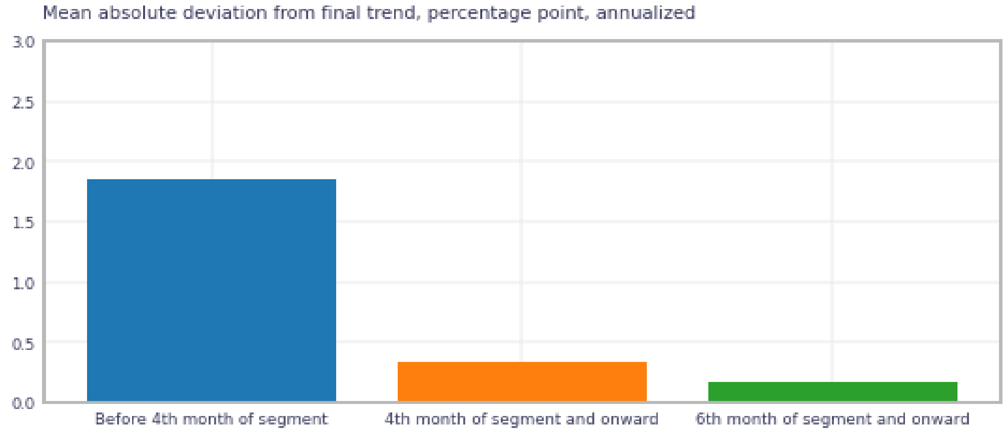

The evolution of trend inflation estimates for a representative case (Segment 18, January 2009 to April 2010) illustrates this property (**Chart 6**). Estimates are volatile during the first four months of the segment but converge quickly thereafter. By August 2009, the estimated trend is within 0.4 percentage points of its final value of 1.5 percent. Overall, although segment confirmation requires additional data, useful signals about changes in trend inflation typically emerge within a few months.

Chart 6: Segment 18 illustrates typical real-time segment behavior

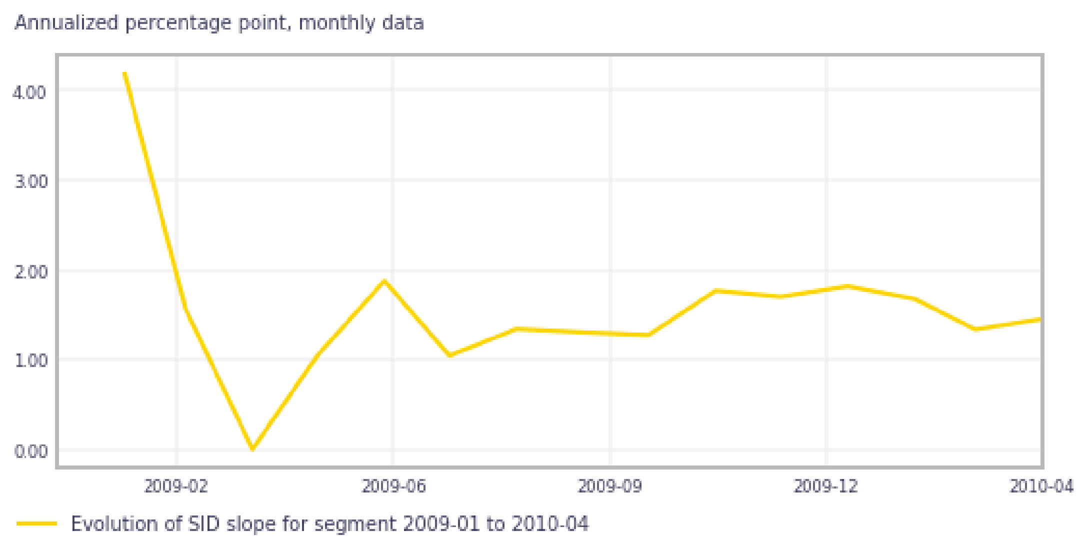

3.4 Current assessment of trend inflation

Using data through December 2025, SID indicates that a new segment began in early 2024 after more than a year with trend inflation near 3 percent (**Chart 7**). The estimated trend initially fell below 2 percent and then returned to about 2 percent in early 2024.

During this period, year-over-year inflation was affected by persistent shelter cost pressures, changes in indirect taxes, and energy price movements—factors that complicated interpretation. Core inflation measures also varied noticeably.

By contrast, SID adjusted more quickly, indicating a decline in trend inflation earlier than other indicators. While core measures remained elevated or volatile for longer, they gradually converged toward the SID estimate as temporary pressures faded. This highlights SID's ability to filter out short-lived or localized fluctuations and provide a clearer signal of persistent inflation pressures.

Overall, the results suggest that trend inflation has been close to the 2 percent target since early 2024, following the adjustment from the post-pandemic inflation surge.

Chart 7: SID's current trend inflation assessment is stable near 2%

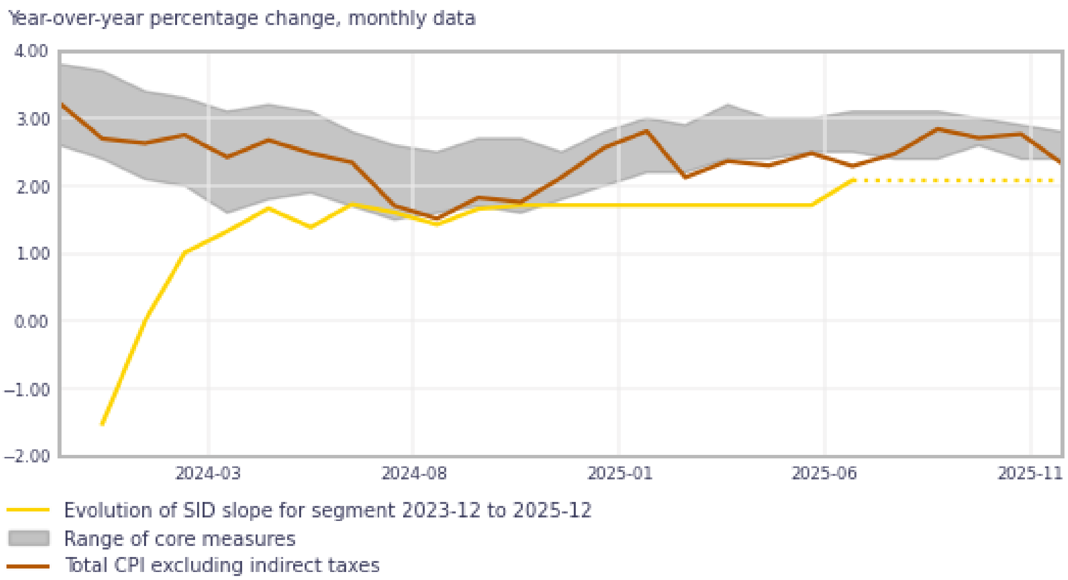

Sources: Statistics Canada and Bank of Canada calculations
Last observation: December 2025

4. Conclusion

This paper introduces Segmented Inflation Dynamics (SID), a simple approach to measuring trend inflation when temporary price movements make underlying inflation pressures difficult to assess. By applying a segmentation algorithm directly to the price index, SID provides a well-defined measure of trend inflation that filters out high-frequency noise and avoids reliance on fixed windows or parametric filters.

The result show that SID responds appropriately to persistent shifts in inflation dynamics. Across specifications, trend estimates are broadly consistent, with current readings clustered near 2 percent. Although estimates early in a segment can be volatile, SID's real-time trend assessment typically converges within four to six months, making it useful when policymakers need timely signals around turning points.

SID is intended to complement existing measures of trend inflation. Its simplicity and transparency make it particularly useful around turning points, when conventional measures may be distorted or slow to adjust. However, SID should not be used in isolation. The method captures persistent changes in inflation dynamics but does not distinguish between broad-based pressures and movements concentrated in a few components. For this reason, SID should be interpreted alongside alternative aggregates and complementary indicators. Used in this way, SID adds to the broader toolkit for distinguishing persistent inflation pressures from temporary price movements, a key challenge for central banks.

References

Armour, J. (2006). "Evaluation of core inflation measures." Bank of Canada Staff Working Paper, 2006-10.

Bank of Canada. (2016). "Renewal of the inflation-control target: Background information." October.

Cheong, S. A., R. P. Fornia, G. H. T. Lee, J. L. Kok, W. S. Yim, D. Y. Xu, and Y. Zhang. (2012). "The Japanese economy in crises: A time series segmentation study." Economics, 6(1).

Dockrill, M., and L. Savoie-Chabot. (2018). "Introducing a systematic measure of idiosyncratic prices." Bank of Canada Staff Analytical Note, 2018-33.

Gordon, R. J. (1975). "Alternative responses of policy to external supply shocks." Brookings Papers on Economic Activity, 6(1), 183–206.

Keogh, E., S. Chu, D. Hart, and M. Pazzani. (2004). "Segmenting time series: A survey and novel approach." In Data Mining in Time Series Databases, edited by M. Last, A. Kandel, and H. Bunke, 1–21. World Scientific.

Khan, M., L. Morel, and P. Sabourin. (2013). "The common component of CPI: An alternative measure of underlying inflation for Canada." Bank of Canada Staff Working Paper, 2013-35.

Khan, M., L. Morel, and P. Sabourin. (2015). "A comprehensive evaluation of measures of core inflation for Canada." Bank of Canada Staff Discussion Paper, 2015-12.

Lao, H., and C. Steyn. (2019). "A comprehensive evaluation of measures of core inflation in Canada: An update." Bank of Canada Staff Discussion Paper, 2019-9.

Liu, X., Z. Lin, and H. Wang. (2008). "Novel online methods for time series segmentation." IEEE Transactions on Knowledge and Data Engineering, 20(12).

Lovrić, M., M. Milanović, and M. Stamenković. (2014). "Algorithmic methods for segmentation of time series." Journal of Contemporary Economic and Business Issues, 1(1), 31–53.

Macklem, T. (2001). "A new measure of core inflation." Bank of Canada Review, Autumn, 3–12.

O'Brien, D. (2018). "The supercore measure of underlying inflation." In Measures of underlying inflation in the euro area (Box 2). ECB Economic Bulletin, 4.

Porqueddu, M. (2018). “The persistent and common component of inflation (PCCI) measure of underlying inflation.” In Measures of underlying inflation in the euro area (Box 3). ECB Economic Bulletin, 4.

Sousa, A. M. Y. R., H. Takayasu, and M. Takayasu. (2020). “Segmentation of time series in up- and down-trends using the epsilon–tau procedure with application to USD/JPY foreign exchange market data.” PLoS ONE, 15(9).

Sullivan, E. (2022). “Examining recent revisions to CPI-common.” Bank of Canada Staff Analytical Note, 2022-15.

Uzeda, L. (2025). “Pulse check: Measuring underlying inflation and its drivers.” Bank of Canada Staff Analytical Note, 2025-29.

Zhang, Z., J. Jiang, and H. Wang. (2007). “A new segmentation algorithm for stock time series based on the PIP approach.” In Proceedings of the International Conference on Wireless Communications, Networking and Mobile Computing, Shanghai, China, 5609–5612.

Appendix A: Illustration of the enhanced feasibility space algorithm

This appendix provides a step-by-step illustration of the segmentation algorithm.

Figure A1: SID's process of assessing a new data point

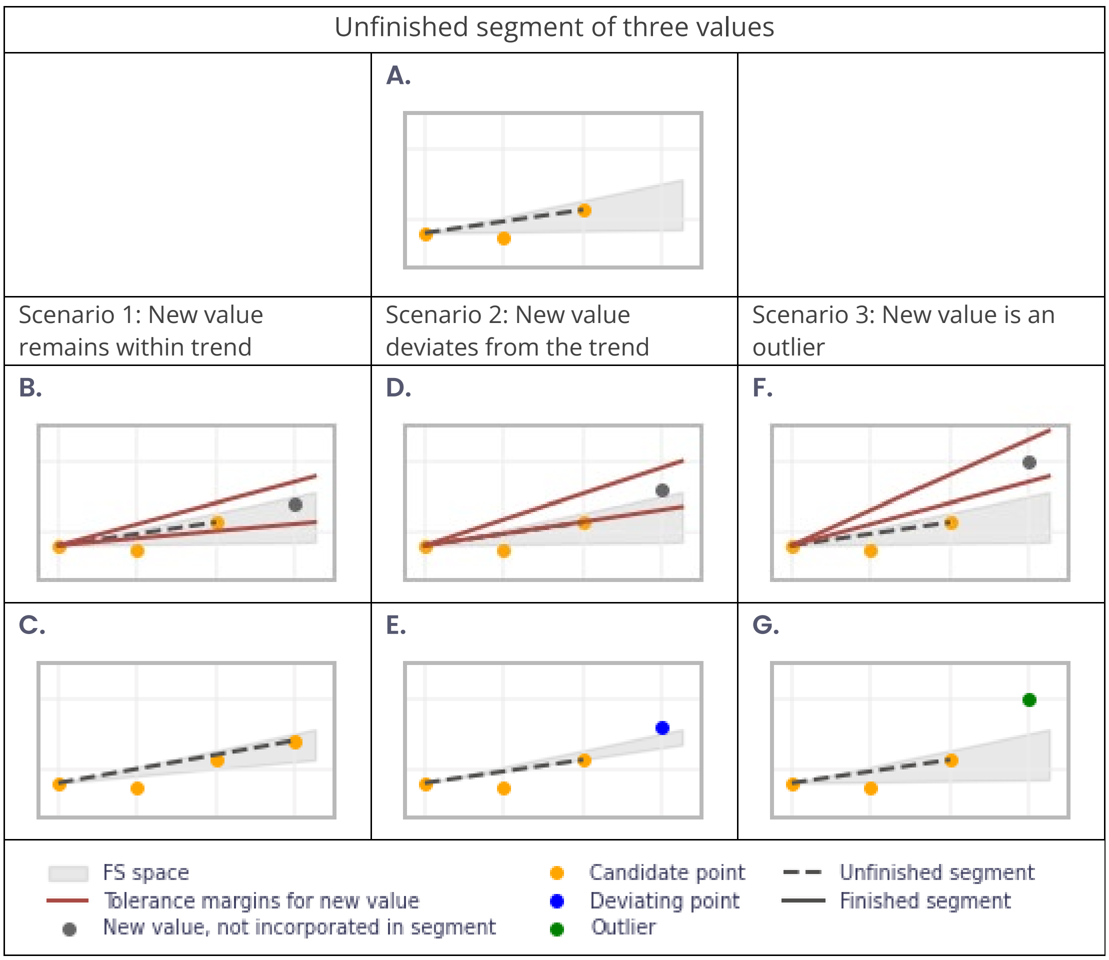

The procedure begins with the first few observations, which establish an initial feasibility space (FS) and an initial interpolation line connecting the first and most recent candidate observation (Figure A1, panel A). The FS includes all lines whose slopes and intercepts place the line within ±δ of every point assigned to the segment.

Updating the feasibility space

When a new observation arrives, tolerance bands of width ±δ are drawn around the observation. The FS is then updated by intersecting the existing FS with the set of admissible lines implied by the new tolerance bands.

The observation is classified as follows:

- **Candidate**: The updated FS is non-empty and intersects with the previous FS. The observation can serve as a potential endpoint, so the interpolation line is extended to this point (Figure A1, panels B and C).
- **Non-candidate**: The updated FS is non-empty but does not intersect with the previous FS. The observation cannot serve as a segment endpoint. The FS is updated, but the interpolation line is not (Figure A1, panels D and E).
- **Outlier**: The updated FS is empty because the tolerance bands around the new observation do not intersect the previous FS. The observation is treated as a temporary disturbance. Neither the FS nor the interpolation line is updated (Figure A1, panels F and G).

In all cases, the algorithm continues as long as the number of consecutive outliers stays below the pre-specified threshold.

Declaring a break

Consider a case in which the consecutive-outlier threshold is three and two consecutive outliers have already occurred (Figure A2, panel A). While this does not trigger a break, it signals a high likelihood of a change in trend inflation.

If the next observation is also classified as an outlier, the threshold is reached and a break is declared (Figure A2, panel B).

Declaring a break has two implications:

- The current segment is closed, and a final interpolation line is drawn between the first and last candidate observations for that segment.
- The procedure restarts from the last candidate observation, which becomes the starting point of the new segment (Figure A2, panel C).

Figure A2: SID ending a segment

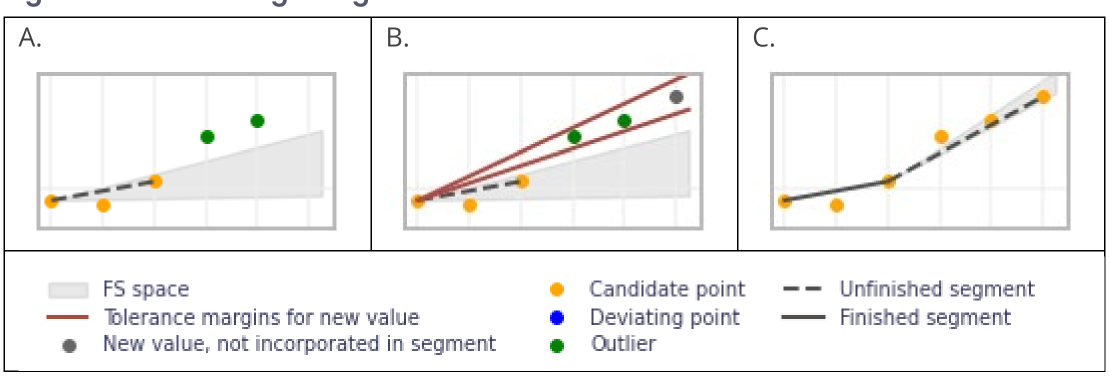

Appendix B: Pseudocode

The following is the pseudocode for SID's process of segmenting a time series.

```
BEGIN
INT max_consecutive_outliers
FLOAT delta
ARRAY[FLOAT] series
INT Index1, Index2 = 0, 1
INT candidate_index = 1
FLOAT pt_slope_new = 0
INT num_consecutive_outliers = 0
ARRAY[ (INT, INT) ] segments = []
FLOAT fs_lower_slope = (series[1] - delta  –series[0]) / (1-0)
FLOAT fs_upper_slope = (series[1] + delta  –series[0]) / (1-0)
###Start segmentation
While index2 <= len(series) – 1:
        ###Calculate slope of tolerance margins for new point
        FLOAT tol_lower_slope = (series[index2] - delta – series[index1]) / (index2 – index1)
        FLOAT tol_upper_slope = (series[index2] + delta – series[index1]) / (index2 – index1)
        ###Calculate slope to new point
        FLOAT pt_slope_new = (series[index2] – series[index1]) / (index2 – index1)
        ###Intersect with current FS space
        FLOAT fs_lower_slope_new = max(fs_lower_slope, error_lower_slope)
        FLOAT fs_upper_slope_new = min(fs_upper_slope, error_upper_slope)
        if fs_upper_slope_new < fs_lower_slope_new: # Pt is outlier
                num_consecutive_outliers += 1
                if num_consecutive_outliers > max_consecutive_outliers: # Max outliers hit
                        segments.append( (index1, candidate_index) )
                        index1 = candidate_index
                        index2 = index1 + 1
                        candidate_index = index2
                        num_consecutive_outliers = 0
                        fs_lower_slope = (series[index2] – delta – series[index1])
                        fs_upper_slope = (series[index2] + delta – series[index1])
                Else: # Ignore pt
                        index2 += 1
        Else if (pt_slope_new >= fs_lower_slope_new) # Pt is candidate
                and (pt_slope_new <= fs_upper_slope_new):
                ###Update candidate, set fs space to new one and advance index2
                candidate_index = index2
                fs_lower_slope = fs_lower_slope_new
                fs_upper_slope = fs_upper_slope_new
                index2 += 1
        Else # Pt is deviating
                ###Set fs space to new one and advance index2
                fs_lower_slope = fs_lower_slope_new
                fs_upper_slope = fs_upper_slope_new
                index2 += 1
END
```

Appendix C: SID applied to total CPI

Figure C1 highlights the sensitivity of inflation segmentation to indirect tax changes. Because tax measures can create discrete changes in consumer prices, SID applied to total CPI can identify trend shifts that reflect policy actions rather than changes in inflation dynamics. For this reason, the analysis focuses on CPI excluding the effect of indirect taxes.

Many of the differences between the two segmentations occur around periods of major tax changes. These include the federal and provincial tobacco tax reductions in 1994, subsequent increases in tobacco excise tax in the early 2000s, the federal GST rate reductions in 2006 and 2008, and the introduction of the harmonized sales tax in Ontario and British Columbia in 2010. Each of these measures had noticeable effects on consumer prices without necessarily signalling a change in persistent inflation pressures.

A similar finding emerges in recent years. Segment breaks in total CPI occur around the reinstatement of British Columbia's provincial sales tax in 2012, provincial fuel tax suspensions and rebates in 2022, the federal GST holiday from late 2024 to early 2025, and the removal of the federal carbon price on fuels in April 2025. Excluding the effect of indirect taxes removes the influence of these policy-induced price-level shifts and provides a cleaner signal of trend inflation.

Figure C1: The exclusion of indirect taxes alters SID's assessment of CPI