GMP: Gap-filled Market Profile Universal Construction for Any Data Points

ConQ Research Team

Continual Quasars
April 25, 2026

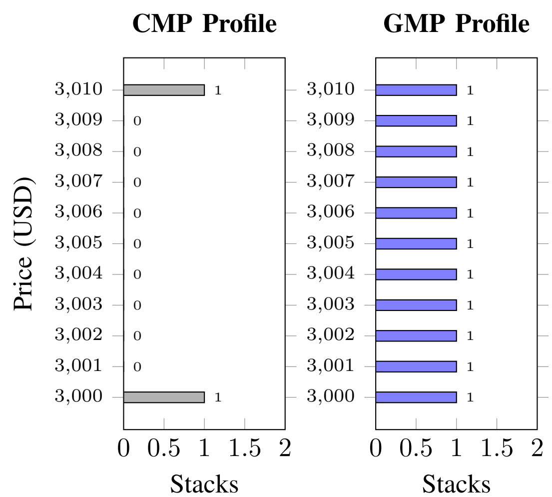

Fig. 1. Horizontal histogram comparison of CMP (left, grey) and GMP (right, blue) for a price move from 3000 to 3010 with $\beta = 1$. CMP shows activity only at the two observed prices; GMP fills all 11 traversed bins. Gap-filling applies to any data sequence.

***Abstract*—Conventional Market Profile (CMP) aggregates price activity into histogram bins, but when applied to any ordered sequence of price points (ticks, candlesticks, or other sampled data), it leaves bins between consecutive points empty. We propose GMP (Gap-Filled Market Profile), a universal construction method that (i) operates on any sequence of price observations and (ii) interpolates every intermediate price bin traversed between successive points, producing a *gap-filled* profile. Building on this gap-filled structure, we introduce an *Up/Down-Bin Footprint Profile* that classifies each bin's contribution directionally, revealing net upward or downward pressure across the price traversal. We formalise CMP and GMP with explicit algorithms, derive the relationship between bin count and a user-defined bin-size parameter β, and present a complete worked example showing how points are grouped into bins under CMP, how gap-filling transforms the sparse CMP output into a dense GMP profile, and how directional footprints are assigned. Charts and tables demonstrate that GMP yields a strictly denser and more informative distribution than CMP, independent of the original data source.**

***Index Terms*—Market Profile, gap-filling interpolation, price bins, directional footprint, high-frequency data, time series**

I. Introduction

The Market Profile, introduced by Steidlmayer [1] and later formalised by Dalton et al. [2], represents price activity as a horizontal histogram whose bins correspond to discrete price levels and whose bar lengths ("stacks") reflect the amount of activity observed at each level. In practice, most implementations construct the profile from candlestick TOCHLV (time, open, close, high, low, volume) data: each candle contributes one stack to every bin between its high and low.

This approach suffers from a fundamental shortcoming when applied to any sequence of discrete price points: when consecutive points are separated by several price levels, the conventional profile records activity only at the two observed prices, ignoring the fact that price must have traversed every intermediate level. This gap neglect is not specific to any data source—it occurs with ticks, candlesticks, or any sampled price series.

We address this issue with **GMP (Gap-Filled Market Profile)**. The construction rule is universal: for every ordered sequence of price observations $\{p_i\}_{i=1}^{N}$, every bin between two successive points receives an interpolated stack, producing a profile with no gaps. The method does not depend on the original data frequency, source, or aggregation level; it applies equally to millisecond tick data, hourly candlesticks, or irregularly sampled price records.

The main contributions of this work are:

1) A rigorous formalisation of CMP and GMP with explicit algorithms and complexity analyses.
2) A theoretical relationship between the user-controlled bin-size parameter β and profile resolution, including a scaling proposition.
3) A universal gap-filling methodology that can be applied to any ordered price sequence, irrespective of source.
4) The introduction of an *Up/Down-Bin Footprint Profile*, a directional classification derived purely from the price traversal without requiring volume or order-book data.
5) A complete, self-contained illustration of the construction on a ten-point price sequence, demonstrating all theoretical constructs.

The remainder of this paper is organised as follows. Section II surveys related work. Section III establishes notation. Section IV defines CMP and GMP formally, presents the GMP algorithm, and introduces the Up/Down-Bin Footprint Profile. Section V provides a complete profile construction example using a 10-point price sequence. Section VI analyses the effect of bin size on profile resolution. Section VII offers a minimal

TABLE I
NOTATION SUMMARY

Symbol	Description
N	Total number of price observations in the sequence
p_i	Price of the i-th observation, $i \in \{1, \ldots, N\}$
β	Bin size (price units per bin); default $\beta = 1$
$b(p)$	Bin index of price p: $b(p) = \lfloor p/\beta \rfloor$
$S[k]$	Stack count (profile value) at bin k
Δ_i	Price displacement: $\Delta_i = p_i - p_{i-1}$
K_i	Number of bins traversed from observation $i-1$ to i
$U[k]$	Up-bin count at bin k
$D[k]$	Down-bin count at bin k
$\delta[k]$	Net footprint delta at bin k: $\delta[k] = U[k] - D[k]$

illustrative example. Section VIII discusses theoretical implications, and Section IX concludes.

II. Related Work

A. Market Profile

The Market Profile concept originates with Steidlmayer's observation that price distributions at each level reveal where market participants find "fair value" [1]. Dalton et al. [2] extended the framework with auction-market theory, using half-hour brackets as time-price opportunities (TPOs). Both formulations rely on time-based bars rather than raw ticks, but the underlying logic of binning price activity is independent of the data source.

B. Interpolation in Financial Time Series

Interpolation techniques are common in high-frequency finance. Clark [3] demonstrated that subordinating returns to trade-count time yields closer-to-Gaussian distributions, motivating trade-indexed analysis. Ané and Geman [4] confirmed that business-time transformations normalise returns at the tick level. The gap-filling approach we propose is conceptually similar to linear interpolation on the price axis, but applied to histogram bin counts rather than to prices themselves.

C. Footprint and Order-Flow Analysis

Market microstructure theory, including Glosten and Milgrom [5], O'Hara [6], and Madhavan [7], provides foundations for analysing directional pressure. Traditional footprint charts distinguish trades at bid versus ask prices. Our Up/Down-Bin Footprint provides a complementary directional classification derived purely from the sequence of price observations, without requiring volume or order-book data.

III. Preliminaries

Table I summarises the notation used throughout.

[Price observation sequence] A *price observation sequence* is an ordered set $\mathcal{P} = \{(t_i, p_i)\}_{i=1}^{N}$ where t_i is an index (time, trade number, or any monotonic identifier) and p_i is the observed price.

[Bin] Given bin size $\beta > 0$, the *bin* for price p is the integer index

$$b(p) = \left\lfloor \frac{p}{\beta} \right\rfloor . \tag{1}$$

Algorithm 1 CMP Construction

Require: Price sequence $\{p_i\}_{i=1}^{N}$, bin size β
Ensure: Profile array $S_{\text{CMP}}[\cdot]$
Initialise $S_{\text{CMP}}[k] \leftarrow 0 \ \forall k$
for $i = 1$ **to** N **do**
  $k \leftarrow \lfloor p_i/\beta \rfloor$
  $S_{\text{CMP}}[k] \leftarrow S_{\text{CMP}}[k] + 1$
end for
return S_{CMP}

All prices p satisfying $k\beta \le p < (k+1)\beta$ map to bin k.

[Market Profile] A *market profile* is a mapping $S : \mathbb{Z} \to \mathbb{N}_0$ where $S[k]$ counts the number of stacks accumulated at bin k.

IV. Methodology

A. Conventional Market Profile (CMP)

CMP records a stack only at the bin of each observed data point:

$$S_{\text{CMP}}[k] = \sum_{i=1}^{N} \mathbf{1}\big[b(p_i) = k\big], \tag{2}$$

where $\mathbf{1}[\cdot]$ is the indicator function. Bins with no observed point receive $S_{\text{CMP}}[k] = 0$.

Complexity. CMP performs exactly N bin-index computations and N increments, giving $\mathcal{O}(N)$ time complexity.

B. Gap-Filled Market Profile (GMP)

GMP augments CMP by filling every *intermediate* bin between two consecutive observations. The construction proceeds in two phases:

1) **CMP placement.** Each observation p_i contributes one stack to its own bin $b(p_i)$, exactly as in CMP.
2) **Gap-filling.** For each consecutive pair (p_{i-1}, p_i) with $i \ge 2$, every bin *strictly between* $b(p_{i-1})$ and $b(p_i)$ (exclusive of both endpoints) receives one additional stack.

Formally, writing $b_i = b(p_i)$:

$$S_{\text{GMP}}[k] = \underbrace{\sum_{i=1}^{N} \mathbf{1}\big[b_i = k\big]}_{S_{\text{CMP}}[k]} + \sum_{i=2}^{N} \sum_{j=\min(b_{i-1}, b_i)+1}^{\max(b_{i-1}, b_i)-1} \mathbf{1}\big[j = k\big]. \tag{3}$$

When $|b_i - b_{i-1}| \le 1$ (adjacent or same bin), the inner sum is empty and no gap-filling occurs. When $|b_i - b_{i-1}| > 1$, the number of gap-filled (intermediate) bins is

$$G_i = \big|b(p_i) - b(p_{i-1})\big| - 1. \tag{4}$$

The total span of bins traversed, inclusive of both endpoints, is $K_i = G_i + 2 = |b_i - b_{i-1}| + 1$.

Complexity. Let $D = \sum_{i=2}^{N} |b(p_i) - b(p_{i-1})|$ denote the cumulative bin displacement. GMP performs $\mathcal{O}(N + D)$ operations. In the degenerate case where all observations share the same bin, $D = 0$ and GMP reduces to CMP. In the worst case, $D = \mathcal{O}(N \cdot \Delta p_{\max}/\beta)$.

Algorithm 2 GMP Construction (Two-Phase)

Require: Price sequence $\{p_i\}_{i=1}^{N}$, bin size β
Ensure: Profile array $S_{\text{GMP}}[\cdot]$
Initialise $S_{\text{GMP}}[k] \leftarrow 0 \ \forall k$
for $i = 1$ **to** N **do**
  {Phase 1: CMP placement}
  $S_{\text{GMP}}[\lfloor p_i/\beta \rfloor] \leftarrow S_{\text{GMP}}[\lfloor p_i/\beta \rfloor] + 1$
end for
for $i = 2$ **to** N **do**
  {Phase 2: gap-fill}
  $k_{\text{from}} \leftarrow \lfloor p_{i-1}/\beta \rfloor$; $k_{\text{to}} \leftarrow \lfloor p_i/\beta \rfloor$
  if $|k_{\text{to}} - k_{\text{from}}| > 1$ **then**
    $d \leftarrow \text{sign}(k_{\text{to}} - k_{\text{from}})$
    for $k = k_{\text{from}} + d$ **to** $k_{\text{to}} - d$ **step** d **do**
      $S_{\text{GMP}}[k] \leftarrow S_{\text{GMP}}[k] + 1$
    end for
  end if
end for
return S_{GMP}

C. GMP as a Universal Construction

The key contribution of GMP is its universality: the gap-filling rule applies to *any* ordered price sequence, regardless of the original data's temporal spacing, source, or aggregation level. This includes:

- Raw tick data (millisecond-resolution bid/ask records)
- Candlestick TOCHLV sequences (using close, high, low, or any representative price)
- Irregularly sampled price points
- Synthetic price paths or simulated data

The only requirement is that the sequence be ordered (by time, trade index, or any monotonic index). The construction makes no assumption about the mechanism that generated the prices; it purely interpolates bin traversals between consecutive observations.

D. Up/Down-Bin Footprint Profile

Building upon the gap-filled structure of GMP, we introduce a directional classification layer termed the *Up/Down-Bin Footprint Profile*. For every consecutive pair (p_{i-1}, p_i), the trajectory is evaluated as upward or downward based on the price difference. The origin bin $b(p_{i-1})$ is assigned no directional credit for this move (it has already been evaluated by prior action). However, every subsequent bin along the traversed path up to and including the destination bin $b(p_i)$ increments its *up-bin* count $U[k]$ if $p_i > p_{i-1}$, or its *down-bin* count $D[k]$ if $p_i \leq p_{i-1}$.

This algorithm traces the same $\mathcal{O}(N + D)$ bins as the GMP phase, maintaining computational efficiency while providing deep structural insight into directional dominance across the price range.

Algorithm 3 Up/Down-Bin Footprint Construction

Require: Price sequence $\{p_i\}_{i=1}^{N}$, bin size β
Ensure: Profile arrays $U[\cdot], D[\cdot], \delta[\cdot]$
Initialise $U[k] \leftarrow 0, D[k] \leftarrow 0 \ \forall k$
for $i = 2$ **to** N **do**
  $k_{\text{from}} \leftarrow \lfloor p_{i-1}/\beta \rfloor$; $k_{\text{to}} \leftarrow \lfloor p_i/\beta \rfloor$
  if $k_{\text{from}} = k_{\text{to}}$ **then**
    if $p_i > p_{i-1}$ **then**
      $U[k_{\text{from}}] \leftarrow U[k_{\text{from}}] + 1$
    else
      $D[k_{\text{from}}] \leftarrow D[k_{\text{from}}] + 1$
    end if
    continue
  end if
  is_up $\leftarrow (k_{\text{to}} > k_{\text{from}})$
  $d \leftarrow \text{sign}(k_{\text{to}} - k_{\text{from}})$
  $k \leftarrow k_{\text{from}} + d$
  while true do
    if is_up **then**
      $U[k] \leftarrow U[k] + 1$
    else
      $D[k] \leftarrow D[k] + 1$
    end if
    if $k = k_{\text{to}}$ **then**
      break
    end if
    $k \leftarrow k + d$
  end while
end for
for all k **do**
  $\delta[k] \leftarrow U[k] - D[k]$
end for
return U, D, δ

V. Profile Construction Walkthrough

We illustrate the construction on a ten-point price sequence. Each observation is a triple (label, x, y) where *label* is an alphabetic identifier, x is the index, and y the price. Table II lists the data.

TABLE II
INPUT OBSERVATIONS (10 POINTS)

Label	Index #	Price (USD)
A	1	3000.914
B	2	3003.837
C	3	3002.432
D	4	3009.892
E	5	3007.698
F	6	3009.176
G	7	3003.381
H	8	3004.283
I	9	3003.512
J	10	3003.012

A. CMP Profile Table

With $\beta = 1$, the bin index is $b(p) = \lfloor p \rfloor$. CMP counts points per bin. Table III shows that bins 2, 6, 7, and 9 are empty—the gaps in the conventional profile.

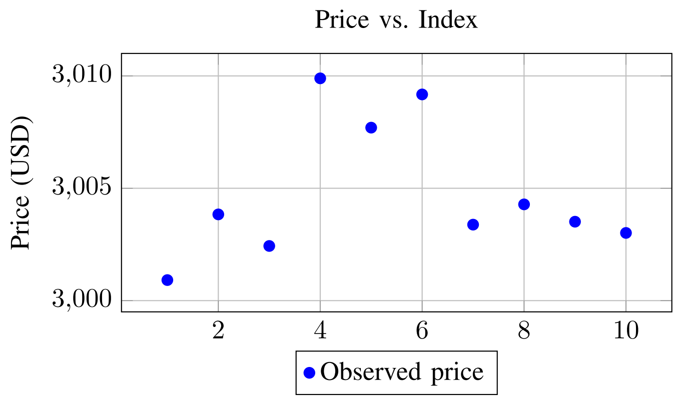

Fig. 2. Price vs. index for the 10-point example.

TABLE III
CMP PROFILE TABLE ($\beta = 1$)

Bin	From	Until	Group	Stacks
1	3000	3001	A	1
2	3001	3002		0
3	3002	3003	C	1
4	3003	3004	BGIJ	4
5	3004	3005	H	1
6	3005	3006		0
7	3006	3007		0
8	3007	3008	E	1
9	3008	3009		0
10	3009	3010	DF	2
		Total stacks		**10**

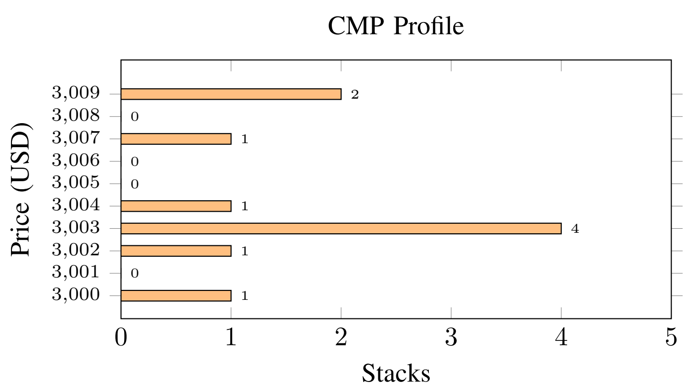

Fig. 3. CMP profile for the 10-point example ($\beta = 1$). Four bins are empty.

B. GMP Profile Table

GMP fills intermediate bins between consecutive points. Table IV shows every bin now populated, with total stack count 25.

C. CMP vs. GMP Side-by-Side

Figure 5 places both profiles side by side. CMP (10 stacks) leaves 40% of bins empty; GMP (25 stacks) fully covers the range.

TABLE IV
GMP PROFILE TABLE ($\beta = 1$)

Bin	From	Until	Group	Stacks
1	3000	3001	A	1
2	3001	3002	A	1
3	3002	3003	AC	2
4	3003	3004	BCGIJ	5
5	3004	3005	CFH	3
6	3005	3006	CF	2
7	3006	3007	CF	2
8	3007	3008	CEF	3
9	3008	3009	CDEF	4
10	3009	3010	DF	2
		Total stacks		**25**

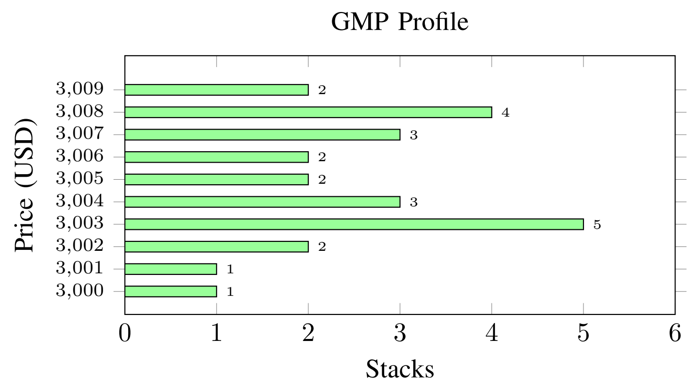

Fig. 4. GMP profile for the 10-point example ($\beta = 1$). Every bin is populated.

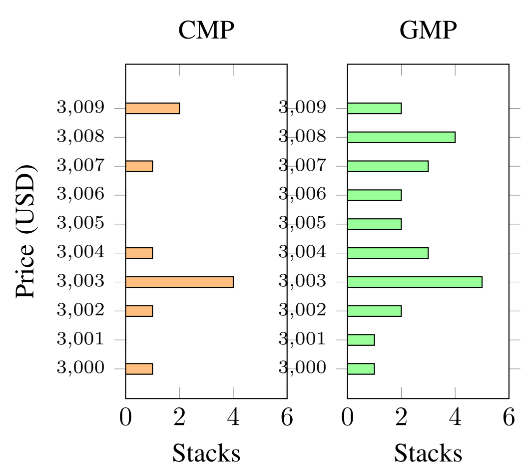

Fig. 5. CMP vs. GMP side-by-side.

Figure 6 presents the entire construction pipeline—raw data, CMP, GMP—in a single three-panel TikZ graphic, ensuring full reproducibility.

D. Up/Down-Bin Footprint Table

Applying Algorithm 3 yields the directional footprint in Table V. For instance, the move from A to B adds up-bins at bins 2–4; the move from C to D adds up-bins at bins 4–10.

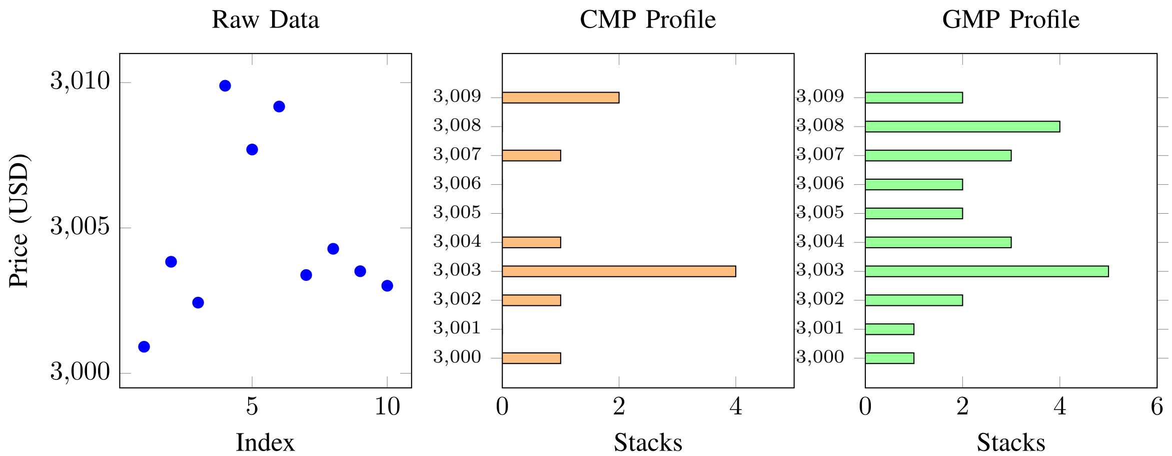

Fig. 6. Three-panel overview: raw data (left), CMP (centre), and GMP (right). Gap-filling produces a continuous profile without empty bins.

TABLE V
UP/DOWN-BIN FOOTPRINT TABLE ($\beta = 1$)

Bin	From	Until	Group	Down	Up	Delta
1	3000	3001	A	0	0	0
2	3001	3002	A	0	1	+1
3	3002	3003	AC	1	1	0
4	3003	3004	BCGIJ	3	2	-1
5	3004	3005	CFH	1	2	+1
6	3005	3006	CF	1	1	0
7	3006	3007	CF	1	1	0
8	3007	3008	CEF	2	1	-1
9	3008	3009	CDEF	2	2	0
10	3009	3010	DF	0	2	+2

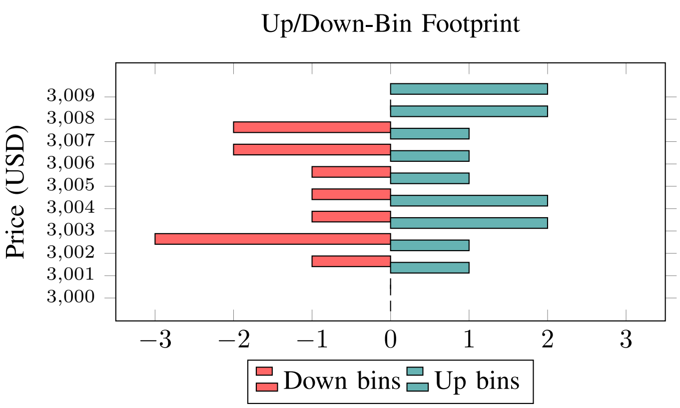

Fig. 7. Directional footprint: down-bins (red) and up-bins (teal).

VI. Effect of Bin Size on Profile Resolution

The bin-size parameter β controls granularity. For a single displacement $\Delta p = |p_i - p_{i-1}|$, the number of traversed bins is

$$K_i(\beta) = \left| \left\lfloor \frac{p_i}{\beta} \right\rfloor - \left\lfloor \frac{p_{i-1}}{\beta} \right\rfloor \right| + 1. \quad (5)$$

Halving β roughly doubles K_i.

TABLE VI
BIN COUNT VS. β FOR $\Delta p = 10$

β	$K_i(\beta)$	CMP bins	GMP bins filled
2.0	6	2	6
1.0	11	2	11
0.5	21	2	21
0.25	41	2	41
0.1	101	2	101

[Bin-count scaling] For fixed Δp and $\beta_1 > \beta_2 > 0$,

$$K_i(\beta_2) \geq \left\lfloor \frac{\beta_1}{\beta_2} \right\rfloor \cdot \big(K_i(\beta_1) - 1\big) + 1. \quad (6)$$

Write $\Delta p = (K_i(\beta_1) - 1)\,\beta_1 + r_1$, $0 \leq r_1 < \beta_1$. Then $K_i(\beta_2) = \lfloor \Delta p/\beta_2 \rfloor + 1 \geq \lfloor (K_i(\beta_1) - 1)\,\beta_1/\beta_2 \rfloor + 1 \geq \lfloor \beta_1/\beta_2 \rfloor\,(K_i(\beta_1) - 1) + 1$.

Table VI quantifies this for $\Delta p = 10$.

Key observations:

- **CMP is invariant:** always 2 bins, regardless of β.
- **GMP scales as $\mathcal{O}(\Delta p/\beta)$:** resolution improves inversely with β. The lower bound is the minimum meaningful price increment.

VII. Illustrative Example

Two observations: $p_1 = 3000$, $p_2 = 3010$, $\beta = 1$. Table VII shows CMP (2 stacks) vs. GMP (11 stacks). Figure 1 plots both.

VIII. Discussion

A. Advantages

1) **No profile gaps.** All traversed price levels are represented, avoiding sparse CMP histograms.
2) **Volume-neutral interpolation.** Interpolated stacks mark traversal, not fabricated volume.
3) **Directional context.** The footprint reveals net pressure without external order-flow data.

TABLE VII
CMP VS. GMP ($\beta = 1$)

Observation #	Price	CMP stacks	GMP stacks
1	3000	1	1
0	3001	0	1
0	3002	0	1
0	3003	0	1
0	3004	0	1
0	3005	0	1
0	3006	0	1
0	3007	0	1
0	3008	0	1
0	3009	0	1
2	3010	1	1
Total stacks		**2**	**11**

4) **Tunable resolution.** β adjusts granularity independently of data frequency.
5) **Universality.** GMP applies to any ordered price sequence.

B. Limitations

1) **Interpolation assumption.** Assumes continuous traversal; genuine price gaps may be over-represented.
2) **Computational cost.** $\mathcal{O}(N + D)$ may be high for large cumulative displacement.
3) **Not a volume profile.** GMP is a pure price-traversal profile; combining with volume is future work.

C. Choosing β

- Set β near the minimum price increment for maximum resolution.
- Enlarge β to reduce noise or align with psychological levels.
- The resolution advantage of GMP over CMP grows as β decreases.

IX. CONCLUSION

We have presented **GMP (Gap-Filled Market Profile)**, a universal construction that interpolates all intermediate price bins between consecutive observations. We formalised CMP and GMP, provided algorithms with complexity analysis, and derived the inverse relationship between bin size β and profile resolution. The detailed 10-point illustration and accompanying charts demonstrate that GMP yields a strictly denser and more informative profile, closing the gaps inherent in conventional methods. Future work may explore weighted interpolation and application across asset classes.

REFERENCES

[1] J. Steidlmayer, *Market Profile*, Chicago Board of Trade, 1986.

[2] J. F. Dalton, E. T. Jones, and R. B. Dalton, *Markets in Profile: Profiting from the Auction Process*, John Wiley & Sons, 2007.

[3] P. K. Clark, "A subordinated stochastic process model with finite variance for speculative prices," *Econometrica*, vol. 41, no. 1, pp. 135–155, 1973.

[4] T. Ané and H. Geman, "Order flow, transaction clock, and normality of asset returns," *The Journal of Finance*, vol. 55, no. 5, pp. 2259–2284, 2000.

[5] L. R. Glosten and P. R. Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, vol. 14, no. 1, pp. 71–100, 1985.

[6] M. O'Hara, *Market Microstructure Theory*, Blackwell, 1995.

[7] A. Madhavan, "Market microstructure: A survey," *Journal of Financial Markets*, vol. 3, no. 3, pp. 205–258, 2000.