

Supplementary Note S1. Mathematical Derivation of the Geometry-Derived Closure Framework

This note derives, from first principles, the system of equations behind the interactive design tool of Section 3.3.3. The goal of the framework is deliberately practical: the user measures a single empirical constant that characterizes how far their membrane deflects, and the tool then returns the geometric dimensions required to print a valve whose inflated membrane conforms to — and therefore seals against — the flow channel. The entire model rests on one exact piece of geometry (a circular arc) plus one experimentally fitted relationship, so it is transparent and easy to reproduce.

We present the linear model first (the default, and the form validated by our own measurements), then generalize it to a nonlinear power law for materials or operating ranges that depart from linearity.

S1.1 Physical system and variables

A microfluidic membrane valve consists of a thin membrane bonded across the top of a flow channel, with a pressurized control chamber above it. When the control chamber is pressurized, the membrane bulges downward into the channel; when it makes full contact with the channel floor (or with a spherical seat recessed into the floor), it occludes the flow and the valve is closed.

We approximate the deflected membrane cross-section as a circular arc and measure all vertical distances downward from the membrane's undeflected *rest line*. This is the same construction shown in main-text Figure 10, in which the flow-channel width W and the membrane width C appear as two parallel chords of a single circle. The model variables are defined in Table 1 of the main text and are used with the same meanings throughout this appendix.

S1.2 The core geometry: two chords of one circle

Consider the arc referred to a vertical axis of symmetry, with depth y measured downward from the rest line. The membrane arc belongs to a circle of radius R whose center O lies on this axis at depth $y = s - R$ (a height $R - s$ above the rest line, since $R > s$). A horizontal line at depth y cuts the circle at a half-width

$$a(y) = \sqrt{R^2 - (y - (s - R))^2} = \sqrt{2R(s - y) - (s - y)^2}. \quad (\text{E1})$$

Two horizontal lines matter. At the *rest line* ($y = 0$) the half-width is $C/2$:

$$(C/2)^2 = 2Rs - s^2 \quad \Rightarrow \quad R = \frac{C^2 + 4s^2}{8s}. \quad (\text{E2})$$

At the *seat top* ($y = g$, where $s - y = s - g = d$) the half-width is $W/2$:

$$(W/2)^2 = 2Rd - d^2 \quad \Rightarrow \quad R = \frac{W^2 + 4d^2}{8d}. \quad (\text{E3})$$

Equations (E2) and (E3) are the *same relation* — the standard chord–sagitta–radius identity — applied to two chords of one circle: the width C at depth 0 and the channel width W at depth g . This is the conceptual heart of the framework: once the membrane is taken to be a circular arc, specifying how it meets the channel walls fixes its entire shape.

S1.3 The closure (sealing) condition

For the valve to seal, the membrane arc must pass through the upper corners of the seat, located at $(\pm W/2, y = g)$. That requirement is exactly Equation (E3): the corner points lie on the circle. Because we *impose* (E3) when we solve the system, every solution the tool returns is a sealing geometry — the design cannot return a leaking configuration. The depth d by which the arc continues below the seat top is then a consequence, not a separate closure test; when H_d is small, the arc bottom simply passes below the floor, which corresponds physically to a spherical seat recessed into the channel floor and carries no “open/closed” meaning.

S1.4 The membrane (constitutive) relation

Geometry alone leaves one degree of freedom: it does not say *how much* a given membrane deflects. That is supplied by a single experimentally fitted relationship between the sagitta s and the membrane width C . For a membrane of fixed material, thickness, and actuation pressure, deflection grows with width, and over a limited range this is well approximated as linear:

$$s = \kappa C. \quad (\text{E4})$$

The constant κ is dimensionless and lumps together the pressure p , elastic modulus E , and thickness t of the specific membrane (Section S1.7 gives the underlying physics). Crucially, κ is measured once on test membranes — it is the only material input the framework needs.

S1.5 Solving the linear system

Given the user’s inputs g, κ, W (and H_d), we solve (E2)–(E4) for the geometry. Substituting $C = s/\kappa$ from (E4) into (E2) gives the radius purely in terms of s and κ :

$$R = \frac{(s/\kappa)^2 + 4s^2}{8s} = \frac{s(1 + 4\kappa^2)}{8\kappa^2}. \quad (\text{E5})$$

Setting (E5) equal to the seat condition (E3) with $d = s - g$:

$$\frac{s(1 + 4\kappa^2)}{8\kappa^2} = \frac{W^2 + 4d^2}{8d} \Rightarrow s(1 + 4\kappa^2)d = \kappa^2(W^2 + 4d^2). \quad (\text{E6})$$

Substituting $d = s - g$ and simplifying (the $4\kappa^2 s$ terms cancel) collapses neatly to

$$(s - g)(s + 4\kappa^2 g) = \kappa^2 W^2, \quad (\text{E7})$$

which is a quadratic in s :

$$s^2 - (1 - 4\kappa^2)gs - \kappa^2(4g^2 + W^2) = 0. \quad (\text{E8})$$

Taking the physical (positive) root and simplifying the discriminant, $(1 - 4\kappa^2)^2 g^2 + 16\kappa^2 g^2 + 4\kappa^2 W^2 = (1 + 4\kappa^2)^2 g^2 + 4\kappa^2 W^2$, gives the closed-form solution used by the tool:

$$s = \frac{(1 - 4\kappa^2)g + \sqrt{(1 + 4\kappa^2)^2 g^2 + 4\kappa^2 W^2}}{2}. \quad (\text{E9})$$

Everything else follows algebraically:

$$C = \frac{s}{\kappa}, \quad R = \frac{s(1 + 4\kappa^2)}{8\kappa^2}, \quad d = s - g, \quad H_c = H_d + g. \quad (\text{E10})$$

A validity requirement is $C > W$: the membrane width must exceed the channel width, otherwise the chord W cannot be a chord of the same circle (equivalently, the seat corners fall outside the arc). The tool flags any input combination that violates this.

S1.6 Obtaining κ from measurements

The design workflow has three steps. *First, measure κ* : fabricate test membranes of several widths C , actuate them at the intended pressure, and measure the resulting sagitta s of each. In practice s need not be measured separately, as it is already contained in the *membrane-reach* data collected for functional characterization (main-text Figure 6). Membrane reach is defined as $\text{reach} = (s/H_c) \times 100\%$, so for a fixed channel height H_c the sagitta is recovered directly as $s = (\text{reach}/100) H_c$. Plotting s against the membrane width C (Figure A1, left) yields a straight line whose slope through the origin is κ . *Second, choose the channel*: fix the flow-channel width W and the desired open lumen g (a fabrication choice; g may be made as small as desired), with the seat height H_d independent and possibly zero. *Third, read off the geometry*: Equations (E9)–(E10) return the membrane width C to print, together with R , s , d , and H_c . Because κ encapsulates the membrane mechanics, the same measured constant designs valves for any channel width in the calibrated range.

A note on units is warranted. The sagitta s is a *vertical* dimension, assembled from print layers (here $50 \mu\text{m}$ each), whereas the width C is a *lateral* dimension, assembled from projector pixels (here $32 \mu\text{m}$ each). Because a layer and a pixel are not the same physical size, s and C must both be converted to a common physical unit (μm) before the ratio $\kappa = s/C$ is formed; the ratio itself is then dimensionless. Working natively in layers and pixels is convenient, as they can be counted directly in microscopy images, provided these nominal conversions are applied consistently. The theoretical layer and pixel sizes used when slicing the STL may differ slightly from the as-printed feature sizes, but this discrepancy is small and *systematic* for a given printer: it is nearly constant across prints made on the same machine and largely cancels in the dimensionless ratio κ . Where absolute accuracy is required, the true pixel and layer sizes can be measured directly against a calibration micro-ruler. The same quantized-unit convention (μm alongside pixels and layers) is used throughout the main text (see Section 2.3).

The fit used to extract κ is illustrated in Figure A1. For NanoClear, this procedure yields $\kappa \approx 0.075$, 0.056 , and 0.037 for membranes of one, two, and three printed layers, respectively ($H = 5$ layers; through-origin fits over the pre-saturation regime): thicker membranes are stiffer and deflect less per unit width, so κ decreases with membrane thickness. Independently, the onset widths below which no deflection is measurable ($W_0 \approx 16$, 44 , and 64 printer pixels, respectively) grow with membrane thickness and set the minimum usable channel width for each membrane recipe.




Figure A1. Determining the membrane deflection coefficient experimentally. Left: for a linear response the sagitta is proportional to the membrane width and the slope is κ . Right: plotting the same data on log–log axes gives a straight line of slope n and intercept $\ln \kappa$; a slope near 1 confirms linearity, while a larger slope indicates the nonlinear model of Section S1.7. Data: measured NanoClear single-membrane-layer ($ML = 1$) deflections from the main-text functional characterization (Figure 6), converted to physical units ($C = W \times 32 \mu\text{m}$; $s = \text{reach} \times H_c$ with $H_c = 250 \mu\text{m}$) over the pre-saturation regime ($W = 40\text{--}80$ printer px, three replicate devices per width; one zero-deflection replicate omitted from the logarithmic panel).

S1.7 Nonlinear generalization: the power-law model

The linear relation (E4) is an empirical convenience, not a law of nature. The mechanics of a clamped, pressurized membrane are genuinely nonlinear, and the deflection scales with width as a *power*:

- *Small-deflection (plate) regime* ($s \ll t$, bending-dominated): the center deflection of a clamped plate under uniform pressure scales as $w \propto p C^4 / (Et^3)$, i.e. $s \propto C^4$, so $n = 4$.
- *Large-deflection (membrane) regime* ($s \gg t$, stretching-dominated): a thin membrane under uniform pressure obeys $s \propto (p/(Et))^{1/3} C^{4/3}$, i.e. $s \propto C^{4/3}$, so $n = 4/3$.

In the case of a compliant resin whose membranes bulge far enough to occlude the channel, deformation is stretching-dominated (the large-deflection membrane regime), which is why $n = 4/3$ is the physically motivated default in the tool. Real devices generally fall between and around these limiting exponents, so rather than hard-wiring $4/3$ we keep n as an adjustable input: the user fits whatever power best collapses their data. Writing the general constitutive law

$$s = \kappa C^n, \quad \Rightarrow \quad C = (s/\kappa)^{1/n}, \quad (\text{E11})$$

and taking logarithms shows how κ and n are obtained simultaneously from one straight-line fit:

$$\ln s = \ln \kappa + n \ln C, \quad (\text{E12})$$

so on log–log axes (Figure A1, right) the slope is n and the intercept is $\ln \kappa$. Setting $n = 1$ recovers the linear model exactly, with κ regaining its original meaning s/C . Note that for $n \neq 1$ the constant $\kappa = s/C^n$ carries n -dependent units, so κ and n must always be reported together. For NanoClear single-layer membranes, the measured exponent over the pre-saturation regime is $n \approx 1.26$ (Figure A1, right) — between the linear ($n = 1$) and large-deflection membrane ($n = 4/3$) limits — so the linear model remains

an adequate engineering approximation, with the power-law option available for materials that deviate further.

Solving the nonlinear system

The geometry (E2)–(E3) is unchanged; only the constitutive law changes. Substituting (E11) into (E2) gives the arc radius as a function of the sagitta,

$$R_{\text{arc}}(s) = \frac{(s/\kappa)^{2/n} + 4s^2}{8s}, \quad (\text{E13})$$

while the seat condition (E3) gives, with $d = s - g$,

$$R_{\text{seal}}(s) = \frac{W^2 + 4(s - g)^2}{8(s - g)}. \quad (\text{E14})$$

The solution is the sagitta at which the two radii coincide, i.e. the root of

$$f(s) \equiv R_{\text{arc}}(s) - R_{\text{seal}}(s) = 0. \quad (\text{E15})$$

For $n \neq 1$ the term $s^{2/n}$ prevents a closed-form solution, so (E15) is solved numerically. The physically meaningful root lies on the branch $g < s < g + W/2$ (equivalently $0 < d < W/2$, a seat contact shallower than a semicircle). On this interval R_{seal} decreases monotonically toward its minimum $W/2$ while R_{arc} increases, so $f(s)$ rises monotonically from $f(g^+) \rightarrow -\infty$ to a positive value; a single root always exists and bisection converges reliably. Once s is found, $C = (s/\kappa)^{1/n}$ and the remaining outputs follow from (E10). As a consistency check, evaluating this numerical scheme at $n = 1$ reproduces the closed-form result (E9) to machine precision.

S1.8 Summary of equations

The closed-form and numerical results for both models are collected in Table A1.

Table A1. Summary of the closure-framework equations for the linear and nonlinear models. Both share the same exact geometry and closure condition; they differ only in the fitted deflection–width relationship.

Quantity	Linear model ($s = \kappa C$)	Power-law model ($s = \kappa C^n$)
Sagitta s	Eq. (E9), closed form	root of $f(s) = 0$, Eq. (E15)
Width C	s/κ	$(s/\kappa)^{1/n}$
Radius R	$s(1 + 4\kappa^2)/(8\kappa^2)$	$[(s/\kappa)^{2/n} + 4s^2]/(8s)$
Intrusion d	$s - g$	$s - g$
Channel height H_c	$H_d + g$	$H_d + g$
Validity	$C > W$	$C > W$

The linear model is the recommended default and matches the response measured for the membranes reported in this work; the power-law model is provided for materials or operating ranges in which the s – C relationship shows curvature.

Supplementary Note S2. Mixed-Pitch (Multirange) Printing Details

Supplementary Note S1 identifies layer thickness as the primary lever for reducing discretization-induced leakage at the seat, but printing an entire device at the finest pitch is counterproductive. At 25 μm the default exposure schedule over-cures negative spaces and partially clogs channels, and membranes formed from two 25 μm layers were over-exposed, burst-prone, and unprintable below a channel width of 90 px. A mixed-pitch (“multirange”) strategy was therefore adopted: the channel and seat regions are printed at 25 μm for maximal seat conformity, while the membrane is formed as a single 50 μm layer. In the multirange test devices the part was printed at 25 μm from the base up to just below the membrane, and at 50 μm from the membrane upward.

Multirange devices initially exhibited a failure mode absent from single-pitch prints: pressurized membranes burst at one edge of the membrane, localized by slow-motion imaging to the boundary between the two layer-thickness ranges. Bursting appeared at channel widths of 100, 110 and 120 px, and only after a sustained pressure hold — not during clearing, filling or the initial application of pressure. That timing rules out a printing or clearing defect and points instead to the local properties of the membrane near the range interface. Single-pitch 50 μm membranes did eventually burst under repeated tests and long holds, whereas multirange membranes burst on the first pressurization. Consistent with a length dependence, short cutout coupons (approximately 3 mm of channel) never burst, whereas full-length channels (approximately 20 mm) did, with or without liquid in the lumen.

Two geometric mitigations were identified. Relocating the range interface at least two 25 μm layers below the membrane plane eliminated bursting at 100 and 110 px, although 120 px still failed after roughly 30 s of sustained hold; filleting the valve-chamber corners then eliminated the edge-burst mode entirely. A confirmation test in which membrane thickness was doubled from 50 to 100 μm also suppressed bursting, supporting a membrane-strength rather than a print-quality origin for the failure.

For anyone pursuing the finer-pitch route, the combination that eliminated the failure was 25 μm seat discretization, a single 50 μm membrane layer, a range interface recessed at least two layers below the membrane plane, and filleted valve-chamber corners. Because the complications of mixed-pitch printing outweighed its benefit for the seat geometries adopted in the main text (see Section 3.4.1), all devices reported there were printed single-pitch; this note records the investigation so that the failure mode and its remedies are available to anyone attempting the finer-pitch route.

Supplementary Note S3. Print Recipe for Reported Devices

1. Single-range prints (all reported devices except the mixed-pitch tests of Section 3.4.2). Asiga Composer (STL slicer) settings:
 - Default parameters from Asiga’s Material Library. These include exposure time, waiting time, build plate velocity, temperature, pressure sensors, etc.
 - Uncheck anti-aliasing option.
 - Uncheck Z-compensation option.
2. Mixed-pitch (multirange) prints (Section 3.4.2):

- Default parameters from Asiga's Material Library. These include exposure time, wait time, build plate velocity, temperature, etc.
- Uncheck anti-aliasing option.
- Uncheck Z-compensation option.
- Modified print range options: instead of print from 0 (mm) to max height (in mm), the prints were split into two distinct ranges: from 0 to a designated height (in this case, (i) right before the membrane, and (ii) 2 layers before the membrane) of a lower layer thickness option (25 μm), and then from designated height to max height in a higher layer thickness option (50 μm).
- Each of the two ranges has a different set of default parameters, which is automatically adjusted by the Asiga Composer (version 2.1.9), and were left as is for the prints.

3. Materials:

- Commercially available resin FunToDo's NanoClear.
- Commercially available resin Asiga's Dental IBT.

4. Printer settings:

- The printer was calibrated for balancing (referred to as "Zero Positioning" on Asiga printers) before each print.
- Enhanced modes, particularly the 4K mode and Transparent mode were disabled for these microfluidics prints.
- Print trays (resin vats) were checked after every 10 prints to ensure they were free of debris, artifacts, or damaged film: all liquid resin was filtered out, the tray film was cleaned and visually inspected, and the tray was refilled with the filtered resin.

5. Composer (STL slicer) settings:

- STL files imported into Asiga Composer were kept in their design orientation (all three X-Y-Z axes parallel to the Composer axes).
- All prints were manually placed at the center of the build plate; the exact placement varied with print size and the number of simultaneous prints.
- No support features were added to any print.

Electronic Supplementary Information (ESI): Supplementary Figures

A Geometry-First Design Framework for 3D-Printed Microfluidic Valves Using New Photopolymer Resins




Figure S1. Cross-sections of Functionality Devices at a channel width of 150 printer pixels (4,800 μm) in the unactuated (open, left) and actuated (closed, right) states. At this width the channel roof pre-sags so severely during fabrication that the open-state lumen is already collapsed: the two states are visually indistinguishable, leaving no open baseline against which membrane reach can be defined. Main-text Figure 6 therefore stops at 140 px for this series.




Figure S2. Functionality: valve-closure metrics for thicker membranes at the baseline channel height ($H = 5$ layers, $250 \mu\text{m}$): (a) two-layer ($ML = 2$) and (b) three-layer ($ML = 3$) membranes. Area obstructed (blue, left axes) and membrane reach (red, right axes) versus channel width; symbols and error bars denote mean $\pm$ one standard error of the mean ($n = 3$ replicate devices). The corresponding κ fits are given in Supplementary Note S1.




Figure S3. Functionality (channel-height generality check): **single-layer membranes (ML = 1) at a doubled channel height (H = 10 layers, 500 μm), tested over 20-100 px.** Area obstructed (blue, left axis) and membrane reach (red, right axis) versus channel width; mean ± SEM (n = 3 replicate devices). Fitted membrane deflection coefficient for this series: $\kappa \approx 0.073$ (fitting procedure in Supplementary Note S1).




Figure S4. Comparison of two strategies for improving membrane-to-channel sealing in a pressure-actuated microfluidic valve with a staircase-approximated semi-circular channel floor (e.g., from layer-by-layer fabrication such as 3D printing). Panel labels state the aspect ratio as width:height (AR); the equivalent height-to-width proportions are given below. (a) Baseline configuration with a wide, shallow channel (AR = 3:1, i.e., height-to-width $\approx 1:3$). The deflected membrane (blue) under positive control pressure (P^+) conforms poorly to the coarsely stepped floor profile, resulting in large leak gaps (yellow) between successive contact barriers (red dots) at the staircase corners. The dashed line indicates the ideal smooth circular profile. (b) Height-to-width ratio raised to $\approx 1:1.5$ by increasing channel height while maintaining the same width. Although the membrane deflects deeper, the wide channel preserves a large arc radius, and the horizontal extent of

each leak gap remains substantial: the number of layer steps increases, but the step width — and thus the gap area — remains comparable per step. (c) Height-to-width ratio raised to $\approx 1:1.5$ by narrowing the channel width while maintaining the same height (and layer thickness) as the baseline; the control-chamber width is matched to the reduced channel width. The narrower arc yields tighter membrane curvature relative to the staircase geometry, significantly reducing the triangular leak-gap area at each contact barrier. For staircase-approximated channels, reducing width is therefore more effective than increasing height for improving membrane seal quality at a given layer resolution.




Figure S5. Geometric demonstration of the three canonical valve architectures (Quake, Plunger, and Squeeze) printed monolithically in Asiga Dental IBT, shown in top view (upper row) and side view (lower row); frames are stills from Videos S1-S6. Blue food dye marks the flow channels. Scale bars: 0.4 mm. Each panel is the first footage frame of the corresponding video, immediately after its title card, in which the dye-filled lumen reveals the as-printed channel geometry before actuation; actuation of each valve is shown in the video itself.

Videos S1-S8 each open with a title card identifying the video and close with a laboratory credit card. All eight video files are archived at Zenodo, DOI: 10.5281/zenodo.22051901 (version-independent). Videos S1-S6. Actuation of the IBT valves of Figure S5: S1/S2 Quake (top/side view), S3/S4 Plunger (top/side view), S5/S6 Squeeze (top/side view). All six play at 2× real time. Files:
 ESI_Video_S1_IBT_Quake_top.mp4, ESI_Video_S2_IBT_Quake_side.mp4,
 ESI_Video_S3_IBT_Plunger_top.mp4, ESI_Video_S4_IBT_Plunger_side.mp4,
 ESI_Video_S5_IBT_Squeeze_top.mp4, ESI_Video_S6_IBT_Squeeze_side.mp4.

Video S7. Multi-way flow routing in the three-way redirection device of Figure 15, shown in two parts. S7.1 demonstrates state-retaining sequential addressing: the three arms are filled in the order blue, green, red, each previously filled arm holding its contents, and the device is then flushed. S7.2 demonstrates single-channel addressing by a fill-flush-fill sequence, in which each arm is filled and flushed in turn so that only one arm

carries dye at a time. Both parts are labeled on screen and play at 2× real time. The video is annotated: a dye-colored arrow marks the arm about to be filled before the dye front advances, the three arms are identified as Arm 1 to Arm 3 with an indicator showing what each currently holds, and captions identify every fill and flush step. The two parts were recorded separately, and the stage is repositioned between takes during the flush steps. File: ESI_Video_S7_Flow_Routing.mp4 (contains S7.1 followed by S7.2).

Video S8. Operation of the three-chamber on-chip peristaltic pump of Figure 16, driven through the six-phase actuation sequence of Section 2.8 at a valve actuation frequency of 5 Hz. The frame is split: the left panel is a transmitted-light close-up of the three pump chambers alongside the six-step valve-state table (green = open, red = closed), and the right panel shows the bench configuration, with the fluid reservoir, the three valve pressure inlets, the pump device and the waste container labeled. Dyed fluid is drawn from the reservoir and discharged into the graduated waste container from which the net flow rate of $20.6 \pm 1.3 \mu\text{L/min}$ was determined. Plays at 10× real time. File: ESI_Video_S8_Peristaltic_Pump.mp4.