

Bang-Bang Temperature Control in Espresso Machines

Sawtooth Curve, Extraction Consistency, and the Comparison with PID Control—An Analysis from a Product Development Perspective

Joerg Michael Schudrowitz (ORCID: [0009-0001-0686-8009](https://orcid.org/0009-0001-0686-8009))

Preprint—August 2026 · doi:[10.5281/zenodo.22080631](https://doi.org/10.5281/zenodo.22080631)

Abstract

Consumer-grade espresso machines and bean-to-cup machines primarily regulate their brew temperature using on-off control, known in technical jargon as “bang-bang.” Commercial and prosumer devices, on the other hand, are increasingly relying on PID control. From a product development perspective, this article explains how on-off control works in the key heating architectures—boiler, thermoblock, classic tube-type flow heater, and thick-film heater—why this results in the characteristic sawtooth curve with hysteresis, overshoot, and measurement dead time, and how PID-controlled temperature regulation differs under load. A patent- and manufacturer-supported case study on BSH’s sensoFlow principle demonstrates how a lack of thermal mass is compensated for by sensors at the inlet and outlet, pulsed heating power, and flow rate measurement. It also highlights the new dependencies that arise in this process: the thermal saturation gap and scaling as a disturbance in the control loop. The analysis is based on published literature and patent sources, not on proprietary measurement data. The discussion covers the implications for extraction consistency, reproducibility, and long-term stability, as well as the underlying cost-benefit considerations. The core thesis is: the choice of control and heating architecture is a rational product decision that is, however, invisible to end consumers—and only its disclosure enables informed purchasing decisions.

Keywords: espresso; temperature control; on-off control; bang-bang control; PID control; thermoblock; thick-film heater; flow heater; disturbance feedforward; extraction consistency

Scope and Limitations

This article takes a conceptual-analytical approach and draws on many years of industry experience: the author has worked for over twenty years in product and strategy roles within the home appliance industry—including on the launch of bean-to-cup machines—and is familiar with heating system architecture from industrial practice. The argumentation is based on this experiential knowledge as well as on published scientific literature, manufacturer documentation, and patents—not on the author’s own series of measurements. Quantitative data from non-peer-reviewed industry sources and marketing-oriented statements by manufacturers are identified as such; internal development specifications of individual manufacturers are not asserted because they cannot be publicly verified.

Phrases such as “from the manufacturer’s perspective,” “as seen in development,” or “what one only learns during development” deliberately adopt the insider’s perspective of product development and highlight passages based on the author’s many years of experience in the development and product strategy of bean-to-cup machines. They represent declared experiential knowledge, not findings from the literature—and, unlike the documented sources from literature, patents, and manufacturers, they cannot be readily verified by the reader.

The body of empirical evidence on temperature dynamics is, on the whole, thinner than one might expect. Batali et al. [1] show for filter coffee that, with fixed extraction and strength, the brew temperature has surprisingly little effect on sensory perception. They themselves note, however, that espresso, with its significantly higher concentrations and different extraction dynamics, is a distinct system—so their findings on filter coffee cannot be readily applied to it. The body of evidence specifically for espresso thus remains incomplete and has been studied far less extensively. This article therefore draws on control engineering principles and empirical experience, rather than on its own sensory data.

Note on translation: This English version was translated from the German original using DeepL and subsequently checked against the German source in an AI-assisted review (Claude, Anthropic). The author verified and edited the final text for technical accuracy and domain-specific terminology and takes full responsibility for its content. The German original is the authoritative version and is available at [doi:10.5281/zenodo.22080711](https://doi.org/10.5281/zenodo.22080711).

1. Why Temperature Control Is an Issue in the First Place

It is well documented that water temperature influences espresso extraction. Andueza et al. [2] investigated brew temperatures of 88, 92, 96, and 98 °C and found significant differences in the physicochemical and sensory properties of the espressos. Temperature is one of the sensory-relevant extraction variables. Salamanca et al. [3] further demonstrated that it is not only the setpoint that matters, but also the *temperature profile* during the shot. And Schmieder et al. [4] used fractionated shots to show that the concentration of the extracted components changes significantly over the shot time—the highest concentrations typically occur at the beginning of the shot. Short-term temperature deviations in the first few seconds of a shot thus fall within the time window during which extraction is most intense.

The stability of temperature control is therefore crucial for the extraction result. It is determined by the control architecture; the price segment alone says little about it. Terminologically, this article distinguishes the *temperature profile*—the actual time course of the water temperature—from *temperature control*, the deliberate shaping of that trajectory by the control system.

2. What Bang-Bang Control Is—and How It Works in the Different Heating Architectures

Bang-bang control, known in control engineering as on-off or two-position control, is the simplest conceivable form of temperature control. The heating system has two states. If the measured temperature falls below a lower switching threshold, the heating element turns on at full power; if it exceeds an upper threshold, it turns off. In between lies the hysteresis—a temperature difference defined by the system—which prevents the switching element (typically a relay or a mechanical bimetallic thermostat) from switching on and off every second and wearing out [5]. It therefore exists for the purpose of component protection, not beverage quality; how wide it is set, however, is a trade-off against the achievable temperature stability (Section 7).

This two-point logic applies to fundamentally different heating principles, which differ primarily in their thermal mass. Thermoblocks, tubular flow heaters, and thick-film heaters are flow heaters in the strict sense; the water is heated only at the moment of use. The boiler, by contrast, holds a static

volume of water at a constant temperature. The entire control strategy will later be determined by this spectrum.

In a classic **single-circuit boiler**, the temperature sensor—an NTC (negative temperature coefficient) sensor or, in simple units, a mechanical thermostat—is located on the outside of the boiler wall. It therefore does not measure the water temperature itself, but rather the metal in its vicinity. Anyone who has ever implemented such systems knows that this sensor position results in two distinct errors that must be clearly distinguished. One is a *static offset*—a fundamental, nearly constant temperature difference between the measured metal and the water at the outlet, which can be compensated for by a fixed calibration offset in the software. The other is a *dynamic dead time*: because heat must first travel through the boiler wall to the sensor, the sensor reports every temperature change with a delay—and this time delay cannot be calibrated away; it limits the achievable control performance. This detail does not appear in any product data sheet. At the same time, the boiler is the ideal partner for on-off control: its mass of water acts as a thermal low-pass filter that effortlessly absorbs the residual heat stored in the heating coil after it is switched off. The sluggish controller and the sluggish plant are a perfect match.

The **thermoblock** and the classic **tubular flow heater** heat the water only at the moment of use. While the thermoblock consists of a solid, heated metal block (usually aluminum with a cast-in or pressed-in heating element), in the tubular flow heater the water flows through a directly heated, thin-walled tube—which has significantly less mass than the block. In both cases, the sensor is mounted on the metal, not in the water path. With the high-mass thermoblock, the block temperature is therefore intentionally set significantly higher than the target temperature: the water flowing through it has only a brief contact time with the block, and only a sufficiently hot metal can transfer enough heat during this time to reach the desired brew temperature at the outlet. Technical sources illustrate this by stating that the block temperature is well above the target of 92 °C [6]. This is precisely what gives rise to the characteristic sensitivity to flow rate: with the correct flow rate, the heating power matches the water volume and the outlet temperature is correct; if the flow rate is too fast, the water becomes too cold; if it is too slow, it becomes too hot [6]—a larger, more massive block, whose temperature is closer to the target temperature, smooths out these fluctuations more effectively. Patent US 8,418,601 B2 by BSH Hausgeräte GmbH [7] openly states the consequence: without sufficient thermal mass, the coffee’s outlet temperature depends heavily on the tank water temperature, the ambient temperature, mains voltage fluctuations, and the manufacturing tolerances of the heating element. A thermoblock enables controlled heating and a constant coffee temperature—but requires precisely the large mass for heat storage that drives up warm-up time and material costs. The stated goal of the patent is to operate a coffee machine satisfactorily even *without* a thermoblock, because the lower-mass flow heater is more cost-effective. The proposed solution replaces mass with electronics: the pump operates in pulses, and its duty cycle is adjusted based on temperature, using sensors located before and after the heater [7]. The trade-off between thermal mass, control complexity, and cost is thus clearly documented in the patent literature. It is at the heart of every heating system decision in the product development process.

The third and most recent heating principle is the **thick-film flow heater**, also known as a thick-film heater (both terms are used synonymously below), and colloquially sometimes referred to as a “foil heater.” Instead of a solid metal block, a substrate—typically stainless steel, electrically insulated by a glass-ceramic layer—carries wafer-thin resistive tracks applied by screen printing; the water flows in a channel or tube directly along the heated surface [8]. In U.S. Patent No. 8,600,223 B2 [9], Nestec

describes such a thick-film flow heater for beverage machines, in which even the temperature sensor—a thermistor—is integrated directly onto the thick film, and a control valve continuously adjusts the flow rate to the heater. The key difference lies once again in the thermal mass. Compared to boilers and thermoblocks, it is drastically reduced—with high power density—the OEM manufacturer Otter/Ferro cites 3 kW on a disc with a diameter of 110 mm as an example—and correspondingly faster thermal response [8]. This results in extremely short warm-up times and enables highly dynamic control. In return, the buffering effect is almost completely lost. A boiler smooths out fluctuations through its water mass, while a thermoblock does so through its metal mass. A thick-film heater has virtually nothing left to smooth out; the outlet temperature follows the heating power and flow rate almost immediately.

Thus, these designs form a spectrum of thermal mass: boiler (high), thermoblock (medium), tubular flow heater (low), and thick-film heater (very low). Physically, thermal mass never reaches “zero”—even the substrate, insulation, and heating layers possess heat capacity—but the thick-film heater is at the very lower end of the spectrum. The lower the mass, the less the system stabilizes itself, and the more responsibility for maintaining a constant temperature shifts from the component’s physics to the control system and its electronics. A slow on-off control system, which reacts only after a switching threshold is crossed, places significantly higher demands on sensor placement and hysteresis width in a low-mass thick-film heater and quickly becomes critical: by the time a slow-responding, external sensor registers the steep temperature rise of the high-power-density heating track, the temperature in the flow channel may already have critically overshoot—to the point of local evaporation of the water, which deprives the heating element of its heat-dissipating medium. The thick-film patent literature itself shows that this is precisely the recognized failure mechanism: it documents thermal fuses integrated directly into the thick-film element and channel geometries with heating power that decreases along the flow direction, which are specifically designed to prevent evaporation in the heater and scale buildup [9]. Anyone who has worked in heating system development knows the rule behind this: the heating principle and control strategy are not a freely combinable set of building blocks; they are mutually dependent. The patents demonstrate this—the BSH patent US 8,418,601 B2 [7] compensates for the lack of mass in the cost-effective flow heater through temperature-dependent, cycled pump operation, while the Nestec patent US 8,600,223 B2 [9] couples the thick-film heater to continuous flow control with directly integrated sensors. The mass saved in the component must be compensated for by the electronics through control performance. The long-term outlook, particularly for the latest design, remains unclear: there is hardly any publicly available, reliable data documenting how scale buildup and electronics affect the temperature stability of thick-film heaters over many years of operation; industry sources also cite the lack of long-term experience as an open issue [10].

3. The Sawtooth Curve: Hysteresis, Setpoint, and Switching Thresholds

When observing the temperature curve of an on-off control system over time, the familiar sawtooth pattern becomes apparent (Figure 1). During the heating phase, the temperature rises until it reaches the upper switching threshold, at which point the heating element shuts off. The temperature then drops again due to heat dissipation into the environment. When the water temperature finally falls below the lower switching threshold, the heater turns on again, and the cycle begins anew [5].

Two effects cause the curve to be wider in practice than the hysteresis would suggest. First, **overshoot**: after the heating element shuts off, stored heat remains in it and continues to flow into the

boiler or block—causing the temperature to rise beyond the upper threshold. Second, the **dead time of the measurement**: the sensor is located some distance from the point of action; by the time a temperature change reaches the sensor, it has long since occurred in the water. Together, these factors make the actual temperature amplitude greater than the nominal hysteresis band. Simple single-circuit and thermoblock systems thus fluctuate by several kelvin, while PID-controlled systems fluctuate by a fraction of that. Reliable, independently collected comparative data on this difference is not publicly documented.

For the user, the sawtooth pattern has a practical consequence: the brew temperature of a specific shot depends on where in the cycle the extraction begins. Those who extract immediately after the heater shuts off brew at the upper end of the band. Those who pull a shot shortly before the heater turns back on will brew at the lower end. The practice of “temperature surfing”—deliberately waiting for or triggering a heating cycle before pulling a shot—which is well-known in the barista community, manually compensates for this randomness. The user thus manually takes on a control task that the machine itself does not perform.




Figure 1: Sawtooth curve of bang-bang control (on-off control). Water temperature over time with setpoint as well as upper and lower switching thresholds (hysteresis band); the hatched areas indicate overshoot and undershoot beyond the thresholds due to the heating element’s after-heating effect and the sensor’s measurement dead time. Lower trace: Switching state of the heater (ON/OFF), synchronized with the switching points. The marked shot start points A (near the maximum) and B (near the minimum) illustrate that identically executed shots can start at different brew temperatures. *Schematic representation based on a simplified thermal model; no measurement data.*

4. The Contrast: PID-Controlled Temperature Regulation

PID (proportional–integral–derivative) control works fundamentally differently. Instead of two fixed switching thresholds, it continuously calculates the deviation between the actual and setpoint temperatures and derives a *modulated* heating power from this. The proportional component responds to the current deviation, the integral component eliminates persistent residual deviations, and the derivative component dampens rapid changes and counteracts overshoot [11]. The power is not adjusted continuously but via fast pulse-width modulation: A semiconductor switch—a triac or

solid-state relay—cycles the heating element at second or sub-second intervals with a variable duty cycle. As long as sufficient thermal mass acts as a low-pass filter, it smooths this cycling into a virtually continuously adjustable power output. In systems with very low thermal mass, this applies only to a limited extent: here, mains-synchronous cycling at half- or full-wave intervals (10 or 20 ms at 50 Hz) can produce measurable microfluctuations in the water temperature, so the cycling frequency must be correspondingly high to keep this ripple small. In any case, the switching frequency is several orders of magnitude higher than that of an on-off control system—made possible only by wear-free semiconductor switches; a mechanical relay would not survive the cycling. The control method is never selected in isolation but always as a package comprising sensors, power electronics, and software.

The difference is most evident not during idle operation, but **under load** (Figure 2). When performing an espresso shot, cold water flows into the brew boiler or through the thermoblock—a massive disturbance. An on-off control system only reacts to this once the temperature has fallen below the lower switching threshold, and even then only in a binary manner. A well-tuned PID control detects the drop early based on the gradient and increases power before the drop fully takes effect. Commercial and high-end prosumer machines combine this with a suitably designed architecture. Breville documents in U.S. Patent 10,383,474 B2 [12] a PID-controlled multi-boiler architecture with specific design parameters—a brew boiler with a volume of 300 ml and a heating power of 600 W, i.e., 2 W per ml of boiler volume—and states that the development goal is to maintain a constant brew temperature across multiple shots *and within a single shot*. The three-stage heating process described there is remarkably consistent: preheating via a heat exchanger in the steam boiler, heating to the setpoint temperature in the PID-controlled brew boiler, and temperature maintenance in a separately heated, likewise PID-controlled brew head [12]. This design, together with the control performance, determines how stable the water temperature remains *during the shot*—that is, the temperature that actually comes into contact with the ground coffee.




Figure 2: Bang-Bang vs. PID for a typical load profile. Brew water temperature curve from a low-mass heater during two espresso shots (highlighted in gray: cold water inflow). The green band illustrates an ideal extraction window of $93 \pm 1^\circ\text{C}$. Top: The on-off control shows the characteristic sawtooth pattern during idle operation as well as cycle-dependent start temperatures for the shots. While Shot 1 happens to start near the setpoint (92.5°C), Shot 2 begins at a low point (90.0°C). In addition, the temperature drops sharply under load and overshoots as it settles back. Bottom: The PID control with simulated disturbance feedforward maintains the setpoint precisely during idle operation and starts both shots reproducibly at the setpoint (93.0°C)—this is achieved by the PID component. The fact that the temperature drop under load remains moderate and the response returns without significant overshoot, on the other hand, is due to the disturbance feedforward; a pure feedback loop can only react once the deviation has become measurable. *Schematic representation based on a simplified thermal model; no measurement data.*

5. Software-Based Compensation: The sensoFlow Case Study

Whether software can compensate for the lack of thermal mass in a consumer-grade flow heater is best demonstrated by the most widely documented example—the BSH heating system, which Siemens markets under the name “sensoFlow System” (also advertised as “Intelligent Heater inside”) in its EQ line of bean-to-cup machines and which operates in an identical configuration in Bosch bean-to-cup machines of the Vero series [13].

The verifiable aspects of this can be divided into two categories. Regarding the product claim, Siemens officially states that the sensoFlow system heats the water to the ideal temperature for espresso— 90 to 95°C —and maintains it at a constant level throughout the entire brewing process [13]—a manufacturer’s claim, not an independently measured value. The technical specifications identify the system as a flow heater with a connected load of approximately $1,500\text{ W}$, as seen in the EQ.6 plus s700 [14]; earlier models such as the EQ.8 had a connected load of $1,600\text{ W}$. The technical level is supported by BSH patent literature. One patent application describes a continuous-flow flow

heater deliberately designed to be lightweight, featuring a flattened, thin-walled tube whose mass is characterized as comparatively low [15]. The German parent application for the patent already discussed [7] provides for temperature sensors both in the water outlet between the heater and the brewing chamber and in the inlet, and adjusts the pump duty cycle *and/or heating power* on this basis [16]. Another BSH application takes the decisive step in terms of control technology: pulsed voltage supply to the heating element by selectively blocking individual mains half-waves—fine-grained power control via semiconductors instead of binary on/off—along with two selectively switchable heating elements for power modulation and a flow sensor whose signal, based on the volume of water delivered, is fed into the control system; temperature measurement at the outlet serves as a direct indicator of the beverage outlet temperature [17]. A third patent application controls the pump and heater based on the *difference* between two temperature sensors located on the heater surface and at the outlet [18].

From a control engineering perspective, this involves more than just classic feedback. A pure feedback loop faces a dead-time problem with a heater of low mass: when the pump starts up and pushes cold tank water into the heater, the outlet sensor does not register the drop in temperature until the cold water has reached it—too late for the first few seconds, which are critical for extraction. This dead time arises not only from the path the water takes but also from the thermal resistance between the heating element, the metal, and the water: the incoming cold water first cools the heated block or tube, and the temperature at the outlet sensor drops only after a delay. A flow signal, on the other hand, reports the disturbance the moment it occurs; even the sensor’s startup dynamics are of negligible importance here, because the control system itself initiates the start of the disturbance—the pump startup—and thus knows about it even before the sensor registers the flow. Based on the mass flow rate, the specific heat capacity of water, and the difference between the tank temperature and the setpoint temperature, the required useful thermal power follows directly from the basic equation $P = \dot{m} \cdot c_p \cdot \Delta T$ [19]; the actual electrical power to be supplied is slightly higher due to heat losses to the environment ($P_{el} = \dot{m} \cdot c_p \cdot \Delta T / \eta$), and in the transient case, the control system must additionally account for the heat capacity of the component itself. The control system can thus increase power *before* the sensor detects the drop—the essence of disturbance feedforward control. The combination of a flow sensor and pulsed heating power documented in [17] provides exactly the sensor and actuator technology required for this. The exact form of the algorithm in the current production firmware is not publicly available; what is documented, however, is the principle—sensors at the inlet and outlet, half-wave-pulsed heating power, and flow measurement—as well as the product’s promise of a constant brew temperature within the target range. Further statements regarding the specific production implementation cannot be publicly verified and are therefore not made here.

The case study is doubly illuminating for the article’s central thesis. On the one hand, it confirms the possibility of compensation: the trade-off between rapid readiness and energy efficiency on the one hand, and temperature stability on the other, cannot be resolved through hardware alone without long warm-up times and standby losses—here, the complexity shifts from the hardware to the sensors and software. On the other hand, it demonstrates the limitations of these categories: from a control engineering perspective, such a system is no longer an on-off control system, even if it is based on a low-cost flow heater. The equation “cheap flow heater = poor temperature control” therefore falls short; what matters is the interplay of thermal mass, sensor technology, and control. The requirement for transparency remains unaffected by this: from the outside, the customer cannot tell whether such a control system is installed, nor how well it is tuned.

6. Practical Implications for Extraction Consistency

The practical implications affect four levels: reproducibility from shot to shot, the temperature profile within a single shot, the warm-up phase, and long-term stability. In all four cases, these are consequences expected based on control theory and the cited literature, not results of the author's own measurements. Section 6.5 concludes by assessing the limits of this evidence.

6.1 Shot-to-Shot Reproducibility

The most directly affected aspect is reproducibility from shot to shot. If the starting temperature varies by several kelvin depending on the cycle phase, the extraction varies as well—Andueza et al. [2] have shown that temperature differences of this magnitude are sensorially relevant. Two nominally identical shots on the same machine can thus taste different, even though the grind size, dose, or tamping have not changed. For the home user, this is frustratingly vague. The source of the error is invisible, so they look for it in themselves.

6.2 Temperature Profile Within the Shot

Within the shot itself, this variation carries even greater weight because, according to the fractionated measurements by Schmieder et al. [4], the highest concentrations typically occur in the early extraction phase—a drop in temperature in the first few seconds affects the most critical phase. Salamanca et al. [3] also show that the temperature gradient during extraction alters the beverage properties even when the average temperature remains the same. An uncontrolled, plummeting temperature curve thus differs from a deliberately controlled profile, even if both deliver “an average of 92 °C” on paper.

6.3 Thermal Saturation Gap and Warm-Up Phase

Added to this is the thermal saturation gap. The temperature sensor measures the heater—not the path the water takes afterward. There is therefore a systematic gap between “ready indicator on” and “machine actually warmed through”: valves, pipes, and, above all, the group head act as thermal absorbers during the first shot as long as they themselves are still cold. The most massive example is the classic E61 group head from the prosumer portafilter world. Its solid brass construction (over 4 kg) is deliberately designed as a thermal buffer and is passively heated by the boiler via a thermosyphon. Non-peer-reviewed technical sources estimate the time to full thermal saturation at about 20–30 minutes; depending on the machine and the source cited, it can take even longer (up to 40 minutes)—significantly longer than the ready indicator suggests [20,21]. However, the same inertia that smooths out fluctuations during operation also significantly prolongs the warm-up phase. Lightweight bean-to-cup machines drastically shorten this phase but are not immune to the issue; even in these machines, the cold brew path cools the first batch of dispensed water. Manufacturers address this with software—Siemens, for example, documents that EQ bean-to-cup machines automatically rinse after being turned on while cold [22]. This is primarily justified on the grounds of hygiene and self-cleaning; the fact that the same flushing process simultaneously preheats the brew path is a control-engineering interpretation, not a manufacturer's statement. The extent of the thermal saturation gap in individual cases depends on the design, mass distribution, and sensor placement and cannot be reliably quantified without measuring the specific machine. The mechanism can be described—and, as in the E61 example, the order of magnitude reported by technical sources. For the core message of this article, one point is crucial: the control performance at the heater merely sets the upper limit of achievable quality; it does not guarantee it. Even an ideal

PID control system that delivers exactly 93 °C at the outlet is of little use as long as cold metal masses in the brew path act as heat sinks: the group head already draws energy from the water on its way to the puck, and the cold portafilter, together with the filter basket, removes further heat from the coffee during extraction—the temperature arriving at the puck is significantly lower. Temperature stability is therefore a necessary but not sufficient condition; it must be supplemented by a system that is thermally heated all the way to the brewing chamber.

6.4 Long-Term Stability and Aging

Aging takes effect over the longest time scale. Mechanical thermostats drift over their service life; manufacturing tolerances of the heating elements—identified as a disturbance in U.S. Patent 8,418,601 B2 [7]—also cause variation across the production run. The most significant form of aging in practical terms is scaling. A layer of scale between the heating element and the water acts as thermal insulation and impairs heat transfer. In a sluggish, on-off-controlled boiler, this manifests itself relatively benignly as a gradual drop in temperature and longer cycling times. In a low-mass, electronically controlled system, however, the following control mechanism is plausible: the outlet sensor reports that the water is too cold; the controller dutifully increases the power; the additional heat cannot dissipate quickly enough through the layer of scale and accumulates in the heating element—causing the component to overheat. Modern implementations deliberately counteract this thermal runaway: not only through the controller’s derivative component—which reacts only to rates of change, not to dangerous absolute states—but also by limiting the maximum heating rate and, separately, through a dedicated overtemperature monitoring system that shuts down the heater completely if necessary. It is precisely this level of protection that is also documented on the hardware side. In any case, only the individual components of this chain can be verified, not the complete feedback loop as a measured sequence: the patent literature on thick-film heaters documents scaling, evaporation (dry-out), and overheating as real hazards and addresses them through design measures such as thermal fuses integrated into the heating layer and channel geometries specifically intended to prevent evaporation and scale buildup [9]. The linking of these individual phenomena into a closed-loop control loop is a control-engineering interpretation derived from them, not an experimentally documented sequence of events. Consequently, manufacturers of lightweight systems enforce descaling via software: Siemens documents for the calc’nClean program of the EQ series that the timing of the descaling cycle is determined based on the set water hardness and usage, and that the unit remains locked if the cycle is not performed until the service program has run [23]. In these architectures, descaling thus primarily serves to protect a system whose temperature control depends entirely on the electronics. A PID control with an integral component inherently and automatically compensates for slow drift—but it is precisely this compensatory capability that initially masks the scaling rather than eliminating it. A rigid on-off control with fixed thresholds cannot even do that.

6.5 Limits of the Evidence

In contrast, formal modeling of espresso extraction, such as that by Lee et al. [24], typically treats water temperature as a constant boundary condition and focuses on flow and grind size effects. The extent to which real-world temperature dynamics quantitatively shift the extraction is empirically less well measured for espresso than the topic’s taken-for-granted status in the community would suggest—the filter coffee findings by Batali et al. [1] cannot simply be applied, in either direction. Schmieder et al. [4] also urge caution: in their experimental setup, flow rate had the strongest effect on the amounts of components in the cup, whereas the observed effects of temperature were moderate.

Temperature stability is thus primarily an argument for consistency; based on the available evidence, it cannot be justified as the sole determinant of quality.

7. What You Won't Find in Any Data Sheet—The Manufacturer's Perspective

Anyone familiar with heating systems only from data sheets knows about wattage, warm-up times, and perhaps a setpoint. Those in product development recognize the pattern behind this: a chain of compromises. The position of the temperature sensor is a compromise between measurement accuracy, ease of installation, and cost—an externally mounted sensor is robust and inexpensive, but it measures the housing instead of the water and adds dead time; an immersion sensor in the water flow measures quickly, but complicates sealing and design under pressure. The hysteresis width is a compromise between temperature stability and switching life. Thermal mass is a compromise between buffering effect, warm-up time, weight, and material costs—exactly the trade-off described in the BSH patent US 8,418,601 B2 [7], which openly cites cost considerations as its motivation. Overriding all of these are the needs of the target group. A device whose primary use is making milk-based beverages at the push of a button is prioritized differently than a machine for the espresso purist who wants to taste differences of two kelvin.

The point that motivates this article is a different one: the end consumer cannot recognize this architectural decision from the outside. No box says “on-off control, fluctuation range $\pm x$ K.” What is visible are the wattage, warm-up time, and price. Price differences between machines often reflect precisely this decision—a low-mass flow heater with simple control is significantly cheaper to manufacture than a high-mass thermoblock or a PID-controlled multi-boiler system; the patent literature directly addresses this cost difference [7]. The case study in Section 5 also shows that a cost-effective heating principle, combined with appropriate sensors and software, can be expanded into a high-quality temperature control system. This does not make things any easier for the customer, as this level of quality is also invisible from the outside. Both are legitimate product decisions. But only an informed customer can weigh how important temperature stability is to them and what they are paying for. Transparency regarding the control and heating architecture is therefore a key aspect of consumer information; this article aims to contribute to that.

From the manufacturer's perspective, it's important to note: a bang-bang control system is not a design flaw. It is the rational response to a set of requirements in which robustness, cost, and rapid response time are prioritized over the absolute repeatability of the brew temperature. It only becomes problematic when marketing promises and control architecture diverge—when “perfect espresso” is promised, even though the temperature control system is not designed for that purpose.

8. Conclusion: When Which Control System Makes Sense

Both systems can be objectively evaluated. **On-off control** scores points for minimal complexity, high control robustness, low cost, and—with low thermal mass—a short response time. Its trade-off is the sawtooth characteristic: cycle-dependent start temperatures, delayed load response, and no drift compensation. For appliances whose usage profile is dominated by milk-based beverages, convenience, and price, this is an acceptable—and often the right—compromise. **PID control** provides temperature stability during idle and under load, reproducible start-up conditions, and inherent drift compensation—at the cost of semiconductor power electronics, tuning effort, and, in most cases, a deliberately larger thermal mass with a longer warm-up time. For espresso-focused

machines, light roasts with a narrow extraction window, and all applications where reproducibility is the product promise, there is virtually no way around using PID control or a comparably capable continuous control method. Several high-quality espresso machines on the market control temperature using quite different continuous strategies; what matters is the stepless, anticipatory power control, not the PID scheme in the textbook sense.

This is all the more true when the heating principle is taken into account. On a high-mass boiler, on-off control can deliver acceptable results because the mass of water smooths out the sawtooth wave; on a low-mass thick-film heater, fast, well-tuned control becomes practically a prerequisite because the component itself no longer provides any buffering. Control strategy and thermal mass are two sides of the same decision. The sensoFlow case study shows where this decision is shifting in modern home appliance development: complexity is shifting from the hardware to sensor technology and software, which compensate in real time for the reduced thermal mass—using the flow signal as an early warning, pulsed heating power as the manipulated variable, and forced descaling as a protective mechanism for a system whose function depends entirely on the electronics.

The evaluation of a temperature control system thus depends on the alignment between the control system, heating architecture, and product promise. The sawtooth curve and the temperature drop on the first shot are physical consequences of the respective design. Understanding them requires that buyers be familiar with the underlying architecture—establishing this transparency is the aim of this article.

This also leads to the next methodological step. This article draws on control theory and patent analysis; the evidence gap mentioned at the outset—the scarcity of empirical data specifically for espresso—thus remains unaddressed. Closing this gap requires measurement series that directly link the temperature and flow dynamics described here to the result in the cup: real-time flow rate and brew water temperature on the one hand, and extraction yield (measured as Total Dissolved Solids, TDS, and as shot weight) on the other. Only such a comparison can quantitatively demonstrate the extent to which the design differences in the heating and control architectures actually affect the physicochemical consistency of the beverage—and reliably link the machine’s control technology to what the drinker tastes.

Methodological Foundations

This article is a conceptual and analytical overview and perspective piece from the standpoint of product development. It is based on published literature, patent and manufacturer sources, as well as—where specifically noted—the author’s experiential knowledge—not on a series of measurements conducted by the author. The control engineering terms and concepts used (PID control, overshoot, disturbance feedforward/feedforward control) follow the standard presentation of control engineering by Lunze [11]; the description of on-off control with hysteresis follows Schulz and Graf [5]; the fundamentals of heat transfer (thermal conduction, convective heat transfer, thermal resistance, thermal capacity) follow Baehr and Stephan [19]. These standard works serve as a methodological foundation, not as evidence for individual product-specific statements.

References

1. Batali ME, Ristenpart WD, Guinard JX. Brew temperature, at fixed brew strength and extraction, has little impact on the sensory profile of drip brew coffee. *Sci Rep.* 2020;10:16450. doi:10.1038/s41598-020-73341-4

2. Andueza S, Maeztu L, Pascual L, Ibáñez C, de Peña MP, Cid C. Influence of extraction temperature on the final quality of espresso coffee. *J Sci Food Agric*. 2003;83(3):240–8. doi:10.1002/jsfa.1304
3. Salamanca CA, Fiol N, González C, Saez M, Villaescusa I. Extraction of espresso coffee by using gradient of temperature. Effect on physicochemical and sensorial characteristics of espresso. *Food Chem*. 2017;214:622–30. doi:10.1016/j.foodchem.2016.07.120
4. Schmieder BKL, Pannusch VB, Vannieuwenhuysen L, Briesen H, Minceva M. Influence of flow rate, particle size, and temperature on espresso extraction kinetics. *Foods*. 2023;12(15):2871. doi:10.3390/foods12152871
5. Schulz G, Graf K. *Regelungstechnik 1: Lineare und nichtlineare Regelung, rechnergestützter Reglerentwurf*. 5th ed. Berlin, Boston: De Gruyter Oldenbourg; 2015. ISBN 978-3-11-041445-5. doi:10.1515/9783110414462
6. KaffeeWiki. Thermoblock [Internet]. As of February 1, 2011 [cited July 22, 2026]. Available at: <https://www.kaffeewiki.de/wiki/Thermoblock> (specialized wiki, not peer-reviewed)
7. Magg J, Mayr A, Obermaier M, inventors; BSH Bosch und Siemens Hausgeräte GmbH, applicant. Device and method for controlling heating processes. U.S. Patent 8,418,601 B2. April 16, 2013.
8. Otter Group. Ferro thick film elements [Internet]. n.d. [cited July 22, 2026]. Available at: <https://ottergroup.co.uk/ferro/thick-film-elements> (OEM industry source, not peer-reviewed)
9. Etter S, Gavillet G, Hodel T, Kollep A, Möri P, Moser R, et al., inventors; Nestec S.A., applicant. Integrated heater for a beverage preparation device. U.S. Patent 8,600,223 B2. 2013.
10. DieRöster. *Espressomaschinen 1x1 Teil 1: Thermoblock, Einkreiser, Zweikreiser, Dual-Boiler* [Internet]. n.d. [cited July 22, 2026]. Available at: <https://www.dieroester.at/kaffeetratsch/espressomaschinen-abc-teil-1-einkreiser-zweikreiser-dual-boiler> (industry source, not peer-reviewed)
11. Lunze J. *Regelungstechnik 1: Systemtheoretische Grundlagen, Analyse und Entwurf einschleifiger Regelungen*. 12th ed. Berlin, Heidelberg: Springer Vieweg; 2020. ISBN 978-3-662-60745-9. doi:10.1007/978-3-662-60746-6
12. Upston G, Corkin DR, Preston A, Grassia R, inventors; Breville Pty Ltd, applicant. Coffee machine with overflow feature. U.S. Patent 10,383,474 B2. August 20, 2019.
13. Siemens Home Appliances (BSH Hausgeräte GmbH). Press release on the EQ.6 bean-to-cup machine—sensoFlow System / Intelligent Heater Inside [Internet]. [cited July 22, 2026]. Available at: <https://www.siemens-home.bsh-group.com/de/siemens-entdecken/presse/pressemitteilungen/detail/pressrelease-1536> (Manufacturer communication)
14. Siemens Home Appliances (BSH Hausgeräte GmbH). Product data sheet for the EQ.6 plus s700 bean-to-cup machine, TE657509DE [Internet]. [cited July 22, 2026]. Available at: <https://media3.bsh-group.com/Documents/specsheet/de-DE/TE657509DE.pdf>
15. BSH Bosch und Siemens Hausgeräte GmbH, applicant. Coffee machine comprising a continuous heater. U.S. Patent Application 2008/0017044 A1. 2008. Available at: <https://patents.google.com/patent/US20080017044/en>
16. Magg J, Mayr A, Obermaier M, inventors; BSH Bosch und Siemens Hausgeräte GmbH, applicant. Verfahren und elektronische Steuereinrichtung zum Steuern von Heizvorgängen in einer Kaffeemaschine. German published patent application DE 10 2004 004 817 A1. 2005. Available at: <https://patents.google.com/patent/DE102004004817A1/de> (German parent application for [7])
17. BSH Bosch und Siemens Hausgeräte GmbH, applicant. Heißgetränkezubereitungsanordnung und Verfahren zum Regeln von Leistungsergebnissen netzspannungsabhängiger Verbraucher in einer Heißgetränkezubereitungsanordnung. German published patent application DE 10 2007 053 253 A1. 2009. Available at: <https://patents.google.com/patent/DE102007053253A1/de>
18. BSH Bosch und Siemens Hausgeräte GmbH, applicant. Temperatursteuerung eines Durchlauferhitzers. German published patent application DE 10 2010 002 440 A1. 2011. Available at: <https://patents.google.com/patent/DE102010002440A1/de>
19. Baehr HD, Stephan K. *Wärme- und Stoffübertragung*. 10th ed. Berlin, Heidelberg: Springer; 2019. ISBN 978-3-662-58440-8. doi:10.1007/978-3-662-58441-5
20. Barista Passione. Die E61 Brühgruppe: Wie eine italienische Erfindung den Brühvorgang beim Espresso revolutionierte [Internet]. September 7, 2019 [updated April 20, 2024; cited July 22, 2026]. Available at: <https://www.barista-passione.de/e61-bruehgruppe/> (coffee blog, not peer-reviewed)
21. Hohlmann B. Die E61 Brühgruppe: Geschichte, Technik und warum sie nach 60 Jahren noch immer dominiert [Internet]. Kaffeemacher; March 4, 2023 [cited July 22, 2026]. Available at: <https://kaffeemacher.de/blogs/kaffeewissen/faema-e61-bruehgruppe> (technical source, not peer-reviewed)
22. Siemens Home Appliances (BSH Hausgeräte GmbH). EQ.500—Daily Maintenance: Cleaning the Brewing Unit (Automatic Rinsing After Power-On) [Internet]. [cited July 22, 2026]. Available at: <https://www.siemens-home.bsh-group.com/de/kundendienst/pflege-reinigung/kaffeemaschinen/eq500/taegliche-pflege/reinigung-brueheinheit>
23. Siemens Home Appliances (BSH Hausgeräte GmbH). Thorough Cleaning: calc'nClean (EQ Series; frequency based on water hardness and usage, appliance lockout if not performed) [Internet]. [cited July 22, 2026]. Available at: <https://>

www.siemens-home.bsh-group.com/at/kundendienst/pflege-wartung/kaffeemaschinen/eq700/gruendliche-reinigung/calc-n-clean

24. Lee WT, Smith A, Arshad A. Uneven extraction in coffee brewing. *Phys Fluids*. 2023;35(5):054110. doi:10.1063/5.0138998 (Preprint: arXiv:2206.12373)

Note on sources: Manufacturer claims regarding temperature stability (e.g., [13]) are product communications, not independent measurements.