VEIRA

CREATIVE VOCAL PLUGIN

VEIRA Vocal Studio

User Manual

An all-in-one creative vocal plugin. A VOICE section shifts pitch and formant, adds unison and tunes the voice; six modules, DYNAMICS, CHARACTER, FILTER, DELAY, REVERB and CHOP, each with one big knob and six styles; and a global bar that levels the voice going in and adds doubling, focus and glue at the end.



Contents

Welcome 3
Getting started 5
What you need 5
Install with VEIRA Manager 5
Activation and the licence 6
Formats and where the files go 6
The panel at a glance 7
Using the controls 8
Quick start: a vocal in five moves 10
How VEIRA Vocal Studio works 11
A voice, four ways to change it 11
What happens inside VEIRA Vocal Studio 11
The latency 12
Controls 13
VOICE 13
The modules 17
DYNAMICS 18
CHARACTER 20
FILTER 21
DELAY 22
REVERB 24
CHOP 25
The global bar 27
Displays 30
The IN and OUT meters 30
Presets 31
The preset browser 31
The preset menu 31
Saving your own presets 31
Presets in your DAW 32
The factory folders 32
Recipes 34
A finished lead vocal 34
Natural pitch correction 34
The hard tune effect 34
A harmony or an octave from one voice 34
A character voice 35
Backing vocals behind the lead 35
A delay throw on one word 35
Stutters and chops on the beat 35
A podcast or voice-over 35
Specifications 36
Troubleshooting and questions 38

Welcome

Thank you for choosing **VEIRA Vocal Studio**. It is one window for everything you might want to do to a recorded voice, from a clean, polished lead to a chipmunk, a robot, a telephone, a wall of harmonies or a stutter on the beat.

The window has three parts:

- **VOICE** reshapes the voice itself: PITCH moves it up or down by up to an octave, FORMANT makes it bigger or smaller without changing the note, UNISON adds detuned copies around it, and CORRECTION pulls it to the notes of a KEY and SCALE, gently or hard. VOICE MIX blends the shifted voice with the natural one: a harmony.
- **Six modules**: DYNAMICS, CHARACTER, FILTER, DELAY, REVERB and CHOP. Each has a power switch, one big knob, one or two small controls and a list of **six styles**: pick a style by name, turn the big knob, and the module does the rest.
- **The global bar** along the bottom: INPUT, AUTO LEVEL, LOW CUT and HIGH CUT going in; MIX with LOCK; DOUBLING, FOCUS EQ, GLUE and OUTPUT going out.

What makes it different

- **Six styles per module instead of a page of knobs.** Every module is a choice of six characters and one amount, tuned so the big knob does what you expect at every setting.
- **The voice itself is a control.** Pitch, formant, unison and correction in one section, with the note sung and the note it goes to shown as you play.
- **It behaves the same on any take.** AUTO LEVEL brings every voice to the same working level before the modules and gives the gain back at the end: a quiet take and a loud one get the same treatment, each at its own level.
- **A mono voice comes out in stereo.** UNISON, DOUBLING, the stereo delays and the reverb spread a mono track across the stereo field; insert it on a mono-to-stereo track.
- **Its own look, easy to read.** The night-sky panel, glowing orbs and glass cards, with every label in clear type, values you can type, and a window you can scale from 75 % to 200 %.
- **80 measured presets** in 13 folders, from leads and backing vocals to tuned, lo-fi, filtered, echoed, chopped and spoken voices. DELAY and CHOP follow the song's tempo.

How it differs from VEIRA Vocal Strip. VEIRA Vocal Strip is a mixing channel strip: nine modules in any order, each with its full set of controls and its own display, for getting a recorded vocal ready for the mix. VEIRA Vocal Studio is the creative tool: a voice shifter with pitch, formant, unison and tuning, and six style-based modules that go from a polished lead to a chopped, crushed, filtered or endlessly reverberant one. Use the strip to mix a vocal; use the studio to make something of it.

> **Note:** VEIRA Vocal Studio has a latency of about 61 ms, most of it the voice shifter's analysis window. It is built for playback and mixing, not for monitoring a singer through it while you record. See *How VEIRA Vocal Studio works ▸ The latency*.

How this manual is organised. *Getting started* installs and activates the plugin, *The panel at a glance* names every part of the window, *How VEIRA Vocal Studio works* shows what happens inside, *Controls* and *Displays* cover every control and readout, with drawings, and *Presets*, *Recipes*, *Specifications* and *Troubleshooting* follow.

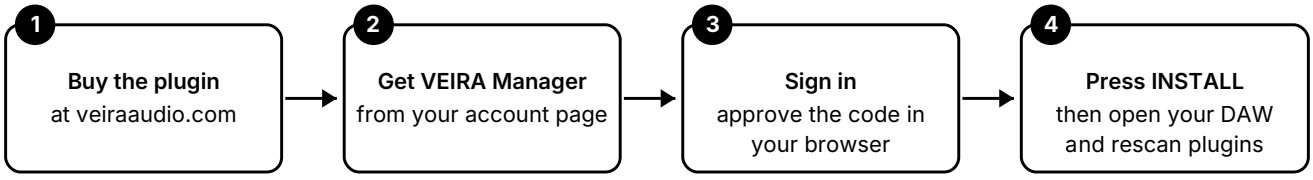
Getting started

What you need

macOS	macOS 11 Big Sur or later. Native on Apple silicon and Intel Macs (universal).
Windows	Windows 10 or 11, 64-bit.
Plugin formats	VST3 (macOS and Windows), Audio Units (macOS), and a standalone app. Pro Tools (AAX) versions arrive once they are signed for Pro Tools.
Hosts	Any DAW that loads VST3 or Audio Units: Logic Pro, Cubase, Nuendo, Ableton Live, Studio One, Reaper, FL Studio, Bitwig and others.
Account	A VEIRA account at veiraaudio.com, and the free **VEIRA Manager** app.
Internet	Needed to install and to refresh the licence. The plugin itself never goes online.

Install with VEIRA Manager

VEIRA plugins are not downloaded from the website one by one. You install them with **VEIRA Manager**, the free VEIRA app that installs, updates and activates everything you own. Its own manual covers it in full; the short version is:



One account works on up to three computers. Updates are free for life.

FIGURE 1 From purchase to first sound in four steps.

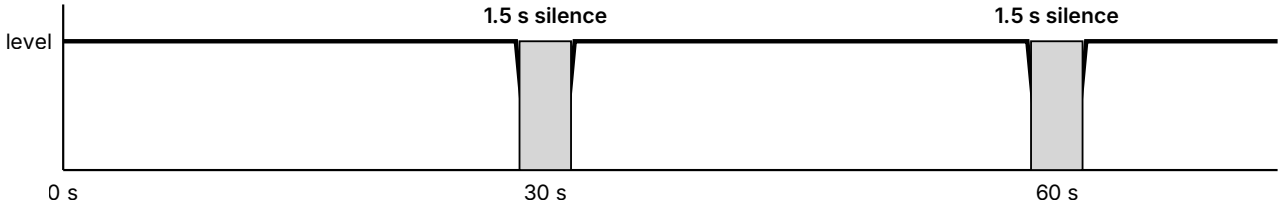
1. **Buy** VEIRA Vocal Studio at **veiraaudio.com**. It's a one-time purchase: you own it for good, with free updates.
2. **Download VEIRA Manager** from your account page (or veiraaudio.com/manager) and install it.
3. **Sign in.** VEIRA Manager shows a short code and opens your browser; confirm the code there. Your password is never typed into the app.
4. Find VEIRA Vocal Studio in the list and press INSTALL. When it says SHOW, it's installed. Open your DAW and let it rescan its plugins.

After installing, VEIRA Vocal Studio appears in your DAW under the manufacturer **VEIRA** (in Logic Pro: Audio Units ▸ VEIRA ▸ VEIRA Vocal Studio).

Activation and the licence

VEIRA Manager keeps a licence on your computer that the plugins check. It is refreshed every time VEIRA Manager opens and stays valid for 30 days offline, so you can work without internet as long as you open VEIRA Manager now and then.

Without a valid licence, VEIRA Vocal Studio still loads and works, but **its sound fades out for 1.5 seconds every 30 seconds**, and a strip along the bottom of its window says why (for example, *NOT ACTIVATED ON THIS COMPUTER. SIGN IN WITH VEIRA MANAGER.*). Clicking the strip opens your account page.



Without a licence: short silences every 30 s, faded in and out so nothing clicks.

FIGURE 2 What an unlicensed copy does. With a valid licence the audio is never touched.

- **Activating while your DAW is open** takes effect within about two minutes. You don't need to reload the plugin.
- **"Not activated on this computer"**: open VEIRA Manager and sign in.
- **"Expired"**: open VEIRA Manager while online; it renews the licence at once.
- **"Not included"**: the signed-in account doesn't own VEIRA Vocal Studio. Check that you're signed in with the account you bought it with.

Formats and where the files go

VEIRA Manager puts every format in the standard place, so you never need to copy files by hand:

Format	macOS	Windows
VST3	`~/Library/Audio/Plug-Ins/VST3` (just for me, the default) or `/Library/Audio/Plug-Ins/VST3` (all users)	`C:\Program Files\Common Files\VST3` (all users, the default) or `%LOCALAPPDATA%\Programs\Common\VST3`
Audio Units	`~/Library/Audio/Plug-Ins/Components` or `/Library/Audio/Plug-Ins/Components`	(Audio Units are macOS only)
Standalone	`~/Applications/VEIRA` or `/Applications/VEIRA`	`C:\Program Files\VEIRA` or `%LOCALAPPDATA%\Programs\VEIRA`

You choose *just for me* or *all users* in VEIRA Manager's install settings.

The standalone app runs VEIRA Vocal Studio on its own, with your audio interface's input and output, which is handy for trying presets without a DAW.

The panel at a glance



FIGURE 3 VEIRA Vocal Studio with the factory preset *Pop: Radio Pop Lead* loaded while a vocal plays: DYNAMICS (*Shine*), DELAY (*Ping-Pong*) and REVERB (*Plate*) are on, CHARACTER, FILTER and CHOP are off.

	Part	What it does
1	Preset browser	Steps through and saves presets. See *Presets*.
2	VOICE switch	The VOICE section on or off. Off, the four orbs and VOICE MIX dim.
3	VOICE MIX	The shifted voice against the natural one.
4	PITCH	Moves the voice up or down, −12 to +12 semitones.
5	LINK	FORMANT follows PITCH: the chipmunk and giant voices.
6	FORMANT	The size of the voice, −12 to +12 semitones, without changing the note.
7	UNISON	Detuned copies around the voice, spread left and right.
8	CORRECTION	Pulls the voice to the notes of KEY and SCALE.
9	KEY and SCALE	The notes CORRECTION may use.
10	NOTE	The note being sung and, with CORRECTION on, the note it goes to.

	Part	What it does
11	IN and OUT meters	The level going in and coming out.
12	Power switch	One per module: the module on or off.
13	The big knob	The module's main control: AMOUNT, CUTOFF or MIX.
14	Small controls	The module's one or two secondary knobs (here GATE, with the GR readout beside it).
15	Style list	Six styles; click one to choose it. The chosen one sits in a lighter bar.
16	PRE	CHOP before DELAY and REVERB, so their tails flow on.
17	INPUT	The level going in.
18	AUTO LEVEL	Brings the voice to the working level; the readout shows its gain.
19	LOW CUT and HIGH CUT	Remove rumble below and hiss above.
20	MIX	The whole effect against the dry voice.
21	LOCK	Presets leave MIX where it is.
22	DOUBLING, FOCUS EQ, GLUE	A second take, presence, and a gentle compressor on the whole result.
23	OUTPUT	The level going out.

The window has four areas: the **header** with the preset browser, the **VOICE** section with its four orbs, the **six module cards** between the IN and OUT meters, and the **global bar** along the bottom. The six cards are laid out alike: once you know one, you know them all.

Using the controls

Every VEIRA plugin is played like a piece of hardware. The same few parts appear on every panel:

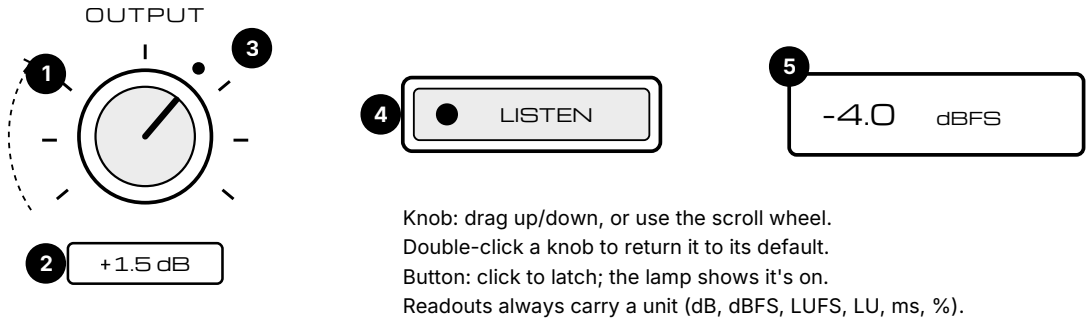


FIGURE 4 The parts you'll find on every VEIRA panel (the labels are examples).

- 1 **Knob.** Drag up to turn it clockwise, down to turn it back, or use your mouse wheel. **Double-click** returns it to its default. The engraved pointer and the small dot on the scale show where it's set.

- (2) **Value window.** The backlit window under each knob shows its exact value with its unit.
- (3) **Silkscreen label.** The control's name, as it appears in your DAW's automation list.
- (4) **Push button.** Click to switch on; the lamp lights and the cap sits lower in its bezel. Click again to switch off.
- (5) **Readout.** Live numbers update ten times a second so you can read them while the meters move. A value over 0 dBFS shows as a warning.

Tooltips: hover over any control or display for a short explanation. On displays, the tooltip depends on where you point.

Automation: every knob and button is a host parameter, so you can automate it in your DAW and control it from a MIDI controller or control surface that your DAW maps.

Units: VEIRA never shows a bare number. **dBFS** is level against digital full scale (0 dBFS is the loudest a digital signal can be); **dB** is a change in gain; **LUFS** is loudness as streaming services and broadcasters measure it; **LU** is a difference in loudness (1 LU = 1 dB).

The glass controls. VEIRA Vocal Studio's panel is dark glass rather than a metal faceplate, and a few things differ from the drawing above:

- **There are no value windows.** The label under a knob (AMOUNT, TIME, OUTPUT…) turns into its value while you hover over the knob or drag it, and back into the label when you leave. The same happens to the MIX caption under the slider.
- **The switches are sliding toggles**, not push buttons: on, the toggle sits to the right and glows (mint on screen; a bright rather than dark toggle on paper). VOICE, the six power switches, LINK, PRE, AUTO LEVEL and LOCK are all toggles.
- **The style lists are menus you can see.** All six styles are printed on the card; click one and it moves into the lighter bar. Hover over a style for a one-line description of it.
- **Type a value:** right-click a knob or an orb (Control-click on a Mac). A field opens with the current value; type the new one and press Return. The number is enough, in the unit the knob shows: **-7** for PITCH, **30** for CORRECTION, **2.5k** or **2500** for FILTER's CUTOFF, **1.6** (seconds) or **800 ms** for REVERB's TIME, **off** for LOW CUT or HIGH CUT. Press Escape or click elsewhere to leave it unchanged. Double-click a knob or orb for its default.
- **Resize the window:** drag its bottom-right corner (the small diagonal lines). It scales from 75 % to 200 % of its normal size and keeps its proportions, so the orbs and the text stay sharp. Your DAW saves the size with the session.

The four orbs (PITCH, FORMANT, UNISON, CORRECTION) are knobs: drag them up or down like any other. Their arc shows the setting: PITCH and FORMANT start from the top and swing left for negative values and right for positive ones; UNISON and CORRECTION fill from the left. The orbs glow with the voice's level while it plays.

Dimmed controls belong to a module that's off, or to the VOICE section when VOICE is off. They keep their values and come back when you switch the module on.

Quick start: a vocal in five moves



Every knob has a short tooltip, and every preset's note says what to turn first.

FIGURE 5 The quickest way to a result.

How VEIRA Vocal Studio works

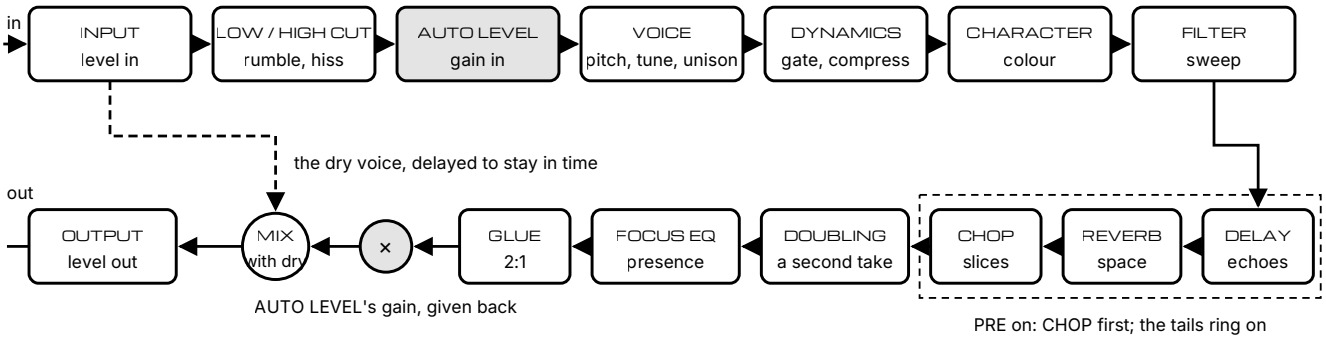
If you use vocal effects every day, skim this chapter; it explains what a voice is made of, what each kind of module does to it, and what happens inside the plugin.

A voice, four ways to change it

A recorded voice has a **note** (its pitch), a **size** (the formants: the resonances of the mouth and throat that make a voice sound like a child, a woman, a man or a giant), a **level** that moves from word to word, a **tone** (dark or bright, clean or broken), and it sits in a **space** and in **time** (the room around it, the echoes behind it, the rhythm of the song). VEIRA Vocal Studio has a section for each:

Part of the voice	What changes it
The note and the size	**VOICE**: PITCH, FORMANT, LINK, UNISON, CORRECTION with KEY and SCALE, VOICE MIX
The level	**DYNAMICS** (six compressors and a gate), AUTO LEVEL, GLUE
The tone	**CHARACTER** (six colours), **FILTER** (six filters), LOW CUT, HIGH CUT, FOCUS EQ
Space and time	**DELAY**, **REVERB**, **CHOP**, DOUBLING

What happens inside VEIRA Vocal Studio



The gain AUTO LEVEL adds before VOICE comes out again (×) after GLUE, so each take leaves at the level it came in.

FIGURE 6 The signal flow. The top row runs left to right, the bottom row right to left. CHOP normally comes after the DELAY and REVERB, so it chops their tails too; with PRE on it moves in front of them.

Step by step:

1. INPUT sets the level going in. A copy is kept, delayed by the plugin's latency, for MIX.
2. LOW CUT **and** HIGH CUT take out rumble under the voice and hiss above it.
3. AUTO LEVEL measures the voice and brings it to a working level of about −18 dBFS, slowly, so that every module that follows sees the same level whatever the take. The gain it used is remembered and taken out again at step 8.

4. **VOICE** shifts the pitch, moves the formants, tunes the voice to the key and adds the unison copies.
5. **DYNAMICS, CHARACTER and FILTER** compress, colour and filter it, in that order.
6. **DELAY, REVERB and CHOP** add the echoes and the room and cut the slices. CHOP comes last, so it chops the echoes and the reverb too; with PRE on it comes first, and the tails ring on over the chopped voice.
7. DOUBLING, FOCUS EQ **and** GLUE add the second take, the presence and the final compression.
8. **AUTO LEVEL's gain is taken back**, in time with the voice, so each take comes out at its own level.
9. MIX blends all of this with the dry voice, and OUTPUT sets the final level.

Every module **always takes its longest latency**, whether it's on or not, so switching a module on or off never shifts the voice; it crossfades over 20 ms instead. Moving CHOP with PRE fades the output out and back in over 15 ms.

The latency

The voice shifter has to hear a little of the voice before it can find its pitch, and CHARACTER's oversampling and the compressors' look-ahead add a little more. Together they make **2,943 samples at 48 kHz (61 ms)**; it is about 61 ms at every sample rate (2,706 samples at 44.1 kHz, 5,855 at 96 kHz, 11,678 at 192 kHz). The latency is reported to your DAW, which moves the track earlier by the same amount on playback, so the voice stays in time with the song. It is **constant**: the same whichever modules are on, bypassed or not, so nothing ever jumps. What it does mean is that a singer can't monitor through the plugin while recording: bypass it while you track and switch it on for playback (see *Troubleshooting*).

Controls

VOICE

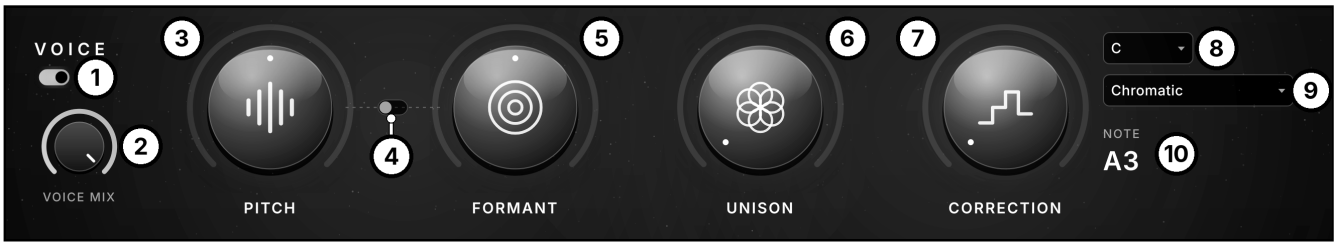


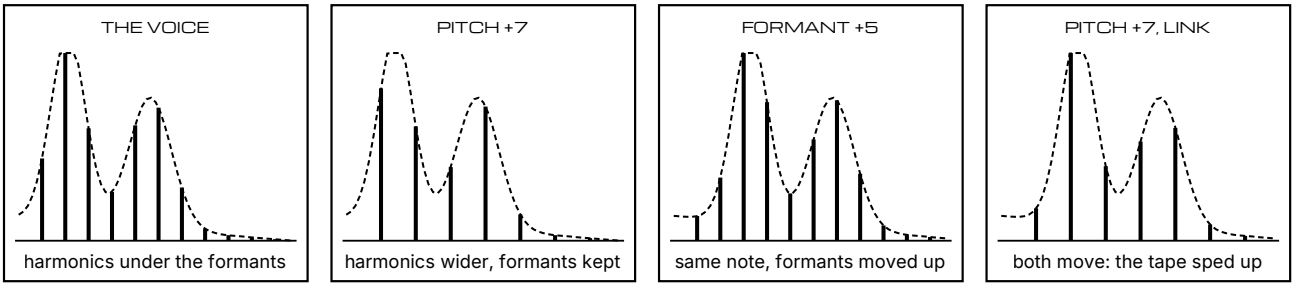
FIGURE 7 The VOICE section (*Pop: Radio Pop Lead*, every orb at its default): 1 the VOICE switch, 2 VOICE MIX, 3 PITCH, 4 LINK with its thread to FORMANT 5, 6 UNISON, 7 CORRECTION, 8 KEY, 9 SCALE and 10 the NOTE readout (here A3 being sung, with CORRECTION at 0 %).

Control	Range	Default	What it does
VOICE	on, off	on	The section on or off. Off, the voice passes through it unchanged (delayed to stay in time).
PITCH	−12 to +12 st in 0.01 st steps	0.00 st	Moves the voice up or down without changing its speed.
FORMANT	−12 to +12 st in 0.01 st steps	0.00 st	Changes the size of the voice without changing the note: down is bigger and darker, up is smaller and brighter.
LINK	on, off	off	FORMANT follows PITCH, like a tape sped up or slowed down. FORMANT's own setting is ignored while on.
UNISON	0 to 100 % in 1 % steps	0 %	Detuned copies of the voice around it, spread left and right.
CORRECTION	0 to 100 % in 1 % steps	0 %	Pulls the voice to the notes of KEY and SCALE: gently at low settings, instantly at 100 %.
KEY	C, C#, D, D#, E, F, F#, G, G#, A, A#, B	C	The key CORRECTION tunes to.
SCALE	Chromatic, Major, Minor, Major Pentatonic, Minor Pentatonic	Chromatic	The notes CORRECTION may use.
VOICE MIX	0 to 100 % in 1 % steps	100 %	The shifted voice against the natural one. 100 % is only the shifted voice; lower blends the two.

How the shifter works. VOICE follows the voice's pitch (it tracks notes from 75 Hz to 1 kHz) and works one period of the voice at a time. To move the pitch it re-spaces those periods closer together or further apart; to move the formants it reads each period faster or slower. Because the two are separate operations, PITCH

and [FORMANT] are independent: you can move one and keep the other exactly where it was. The shifter works on the centre of a stereo track (the mid); the sides pass through, delayed to stay in time. Both knobs glide over 30 ms, so you can automate them without clicks.

PITCH and FORMANT: the note and the size



Left to right in each box: low to high frequency. The bars are the harmonics (the note and its multiples); the dashed line is the formant envelope, the shape of the mouth and throat that makes a voice sound big or small. PITCH moves the bars, FORMANT the line.

FIGURE 8 What the two orbs do to a voice. PITCH moves the harmonics and keeps the formant envelope; FORMANT moves the envelope and keeps the harmonics; LINK moves both together.

A sung note is a row of harmonics, the note and its multiples, under a **formant envelope**: a few broad bumps set by the shape of the singer's mouth and throat. The envelope is what tells you how big the singer is; the spacing of the harmonics is what tells you the note.

- [PITCH] moves the harmonics and leaves the envelope alone, so the same singer sings a different note. **+12 st** is an octave up, **−12 st** an octave down, **+7 st** a fifth. On a sung A (220 Hz), +7 st gives 328.8 Hz, within 0.4 % of the fifth. For harmonies, set [PITCH] to the interval and bring [VOICE MIX] down to 35 to 50 % so the natural voice stays on top (the presets *Fifth Harmony*, *Octave Up Sparkle* and *Octave Down Double* are set up this way).
- [FORMANT] moves the envelope and leaves the harmonics alone, so the same note is sung by a bigger or smaller singer. Down is bigger and darker, up is smaller and brighter. ±6 st moves the centre of a voice's timbre by about a third of an octave each way (measured: from 775 Hz to 545 and 1,080 Hz) with the note unchanged. ±2 to ±4 st changes the character of a singer; ±8 to ±12 st makes a creature.
- [LINK] ties the two: [FORMANT] follows [PITCH], as if the recording were played faster or slower. [PITCH] up with [LINK] on is the chipmunk; [PITCH] down with [LINK] on is the slowed, screwed voice. The dashed thread between the two orbs brightens (mint on screen) while [LINK] is on, and [FORMANT]'s own value is ignored until you switch it off.

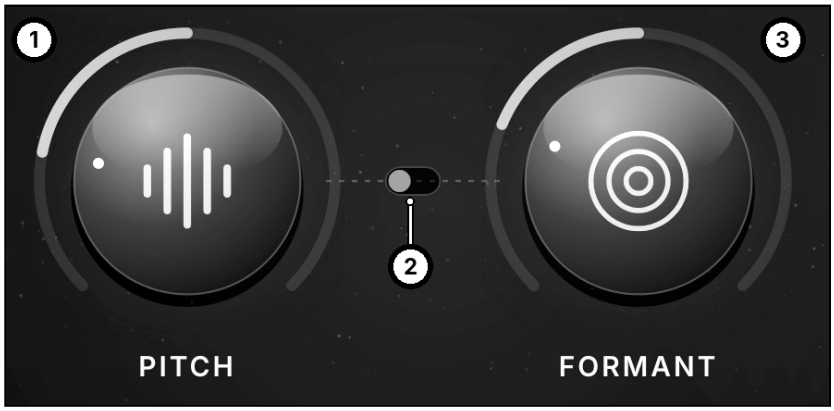


FIGURE 9 *Pitch & Formant: Monster Voice*: 1 PITCH at −7 st and 3 FORMANT at −6 st. The arcs swing left of the top for negative values. 2 LINK is off: the two are set separately.

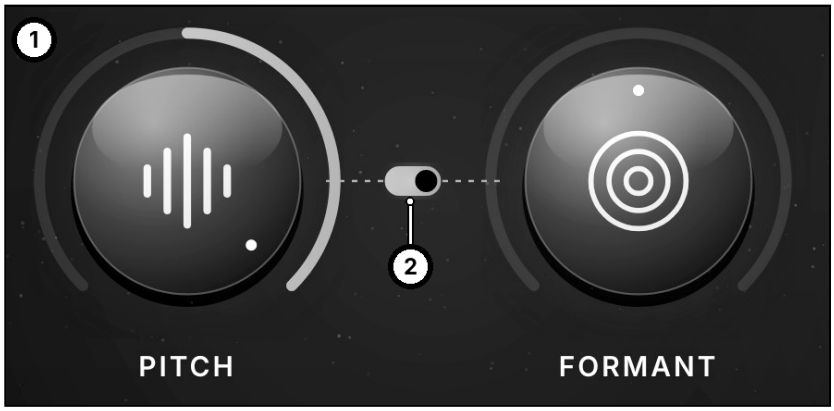


FIGURE 10 *Pitch & Formant: Chipmunk*: 1 PITCH at +12 st with 2 LINK on: the thread between the orbs is bright, and FORMANT follows PITCH whatever its own setting.

> **Tip:** Don't push the pitch and the formant in the same direction unless you want a cartoon. A natural harmony keeps FORMANT at 0; a natural change of character keeps PITCH within ±3 st.

UNISON: a group around the voice

UNISON adds up to four detuned copies of the voice, each a little later than the voice (12, 17, 24 and 31 ms) so that the voice stays in front, and spread left and right so that a mono voice becomes wide. Up to 50 % it adds one pair, panned wide, detuned by up to ±9 cents; from 50 to 100 % a second pair joins nearer the middle, and the detuning grows to ±14 and ±8 cents. The copies are made from the shifted voice, so UNISON on a harmony widens the harmony.

- **15 to 30 %**: a little width and shimmer on a lead (*Wide Chorus Vocal*, *Glassy Pop Tune*).
- **40 to 60 %**: a stacked backing part from one take (*Octave Stack*, *Wide Stack*, *Synth Voice*).
- **80 to 100 %**: a crowd (*Gang Vocals*, *Airy Background Pad*).

On a **mono track** the copies are summed to the middle and UNISON thickens without widening. Insert the plugin on a mono-to-stereo track to hear them spread (see *Troubleshooting*).

CORRECTION: pulling the voice to the key

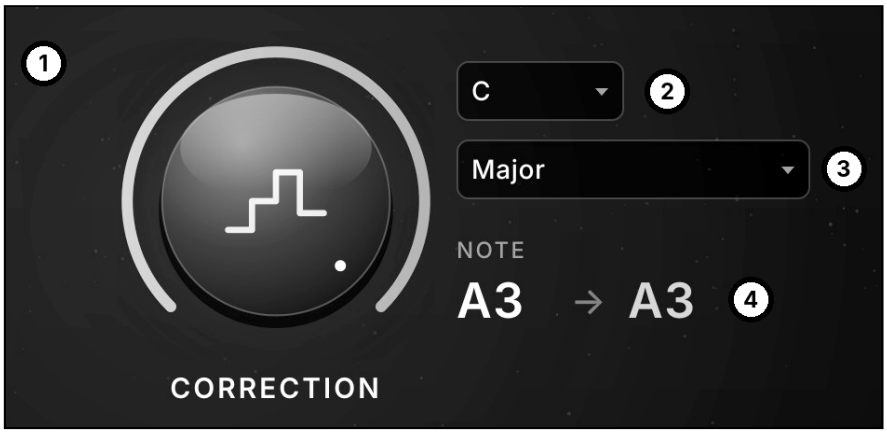
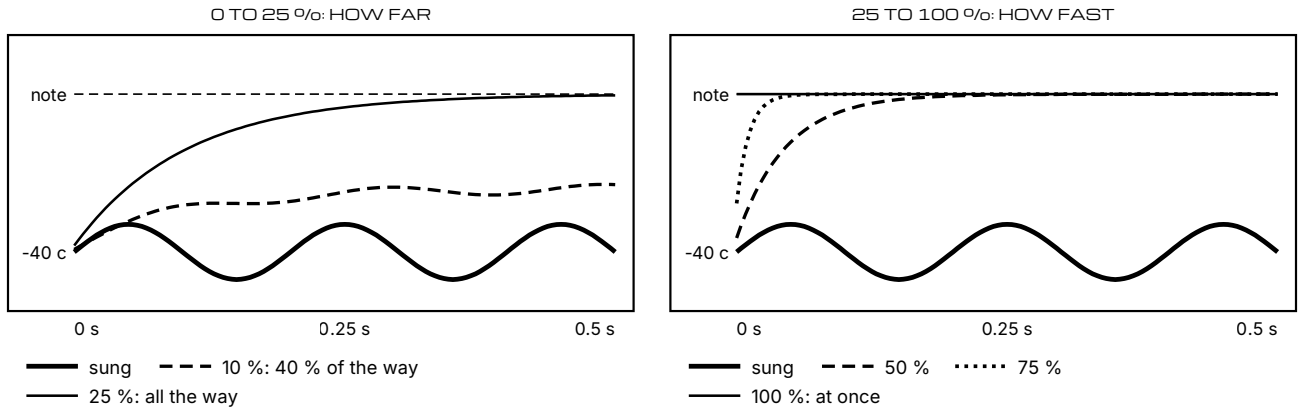


FIGURE 11 *Tuned & Robotic: Hard Tune Major*: 1 CORRECTION at 100 %, 2 KEY C, 3 SCALE Major, and 4 the NOTE readout: the note sung on the left, an arrow, and the note CORRECTION takes it to (mint on screen, the right-hand one on paper).

CORRECTION listens to the note being sung, finds the nearest note of KEY and SCALE, and moves the voice to it. One knob sets both how far and how fast:



A note sung 40 cents flat, with vibrato. Up to 25 %, CORRECTION takes more of the way to the note; above it, the glide shortens from 225 ms at 25 % to instant at 100 %. The vibrato is pulled in with the note: the faster the glide, the flatter the line.

FIGURE 12 CORRECTION's pull on a note sung 40 cents flat. Left: from 0 to 25 % it takes more of the way to the note. Right: above 25 % it goes all the way, faster and faster.

- **0 to 25 %: how far.** At 10 % the voice goes 40 % of the way to the note; at 25 % it goes all the way. The glide is slow here (225 ms at 25 %), so slides and vibrato stay, only tidier.
- **25 to 100 %: how fast.** From 25 % up, CORRECTION always goes all the way, and the glide to the note shortens: 100 ms at 50 %, 25 ms at 75 %, **instant at 100 %.** The faster the glide, the less of the singer's own movement survives: at 100 % every note snaps into place, which is the hard-tuned, robotic effect. A note 40 cents flat lands within 2 cents of the target at 100 %.
- **Natural correction** is 25 to 45 % (*Natural Tune* 30 %, *Tuned Pop Lead* 40 %); **a clearly tuned pop sound** 60 to 80 % (*Trap Vocal* 70 %, *Glassy Pop Tune* 75 %); **the effect** 100 % (*Hard Tune Major*, *Robot Voice*).

KEY **and** SCALE decide which notes CORRECTION may land on. **Chromatic** allows every semitone, so each note goes to the nearest one: safest for melodies with notes outside the key, and the sound of the classic hard tune. **Major** and **Minor** allow the seven notes of the key, and the **Pentatonic** scales only five, so a slightly off note can't land on a neighbour that isn't in the song. The factory presets that use CORRECTION are set to **example keys**: set KEY and SCALE to your song's first, or the voice is pulled to the wrong notes.

The NOTE **readout** shows the note being sung (A3, C#4…) and, as soon as CORRECTION is above 0 %, an arrow and the note it goes to. Watch it while you choose the scale: if the right-hand note is often a semitone away from where the melody wants to be, the key or scale is wrong.

> **Tip:** CORRECTION works on the mid of a stereo track and needs one voice to follow. On a stack of voices or a doubled track it has nothing to lock to; tune each voice on its own track.

VOICE MIX

VOICE MIX blends the shifted voice with the natural one, kept in time with it. At 100 % you hear only the shifted voice (the normal setting for correction and for character voices); lower it and the natural voice comes back underneath, which turns a PITCH setting into a **harmony** and a FORMANT setting into a **thickening**. UNISON's copies are added after the blend, so they stay at their level whatever VOICE MIX says.

The modules

The six cards between the meters work alike:

- **The power switch** at the top left turns the module on or off. Off, the module's controls dim and the voice passes through its place unchanged. Switching crossfades over 20 ms, and the latency doesn't change, so you can switch while the song plays.
- **The big knob** is the module's main control: AMOUNT on DYNAMICS, CHARACTER and CHOP, CUTOFF on FILTER, MIX on DELAY and REVERB. CHARACTER and CHOP at AMOUNT 0 %, and DELAY and REVERB at MIX 0 %, leave the voice exactly as it came in, so you can turn an effect in from nothing. DYNAMICS at 0 % still catches the peaks of the loudest words, and FILTER always filters at its CUTOFF.
- **The small knobs** under it are the module's one or two secondary controls.
- **The style list** at the bottom holds six styles; click one. A change of style switches smoothly, so you can audition the six while the voice plays.

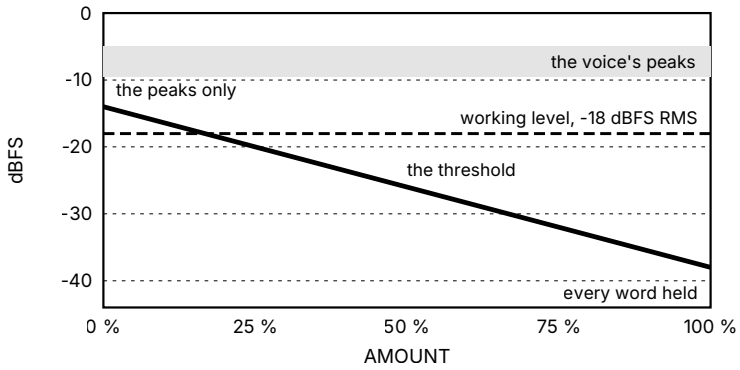
DYNAMICS



FIGURE 13 The DYNAMICS card (*Podcast & Voice: Podcast Host*, *Levelling* 50 %, GATE 30 %) beside 1 the IN meter: 2 the power switch, 3 AMOUNT, 4 GATE, 5 the GR readout (here 10.0 dB of reduction) and 6 the style list.

Control	Range	Default	What it does
Power	on, off	on	The module on or off.
AMOUNT	0 to 100 % in 1 % steps	50 %	How much the voice is compressed: it lowers the threshold.
GATE	0 to 100 % in 1 % steps	0 %	Quiets the gaps between phrases. 0 % is off.
Style	Gentle, Levelling, Solid, Shine, Punchy, Overcooked	Gentle	Which compressor.

DYNAMICS is six compressors in one card, each with a character, and a gate in front of them.



The six styles (ratio, attack / release)

GENTLE
2:1, 20 / 250 ms, soft knee: just holds the voice

LEVELLING
3:1, 30 / 500 ms: evens out phrases

SOLID
4:1, 5 / 120 ms: firm and upfront

SHINE
3:1, 10 / 150 ms, plus air over 8 kHz

PUNCHY
4:1, 35 / 70 ms: the consonants get through

OVERCOOKED
20:1, 0.5 / 50 ms, driven, 70 % wet: smashed

AMOUNT lowers the threshold from -14 dBFS (only the peaks) to -38 dBFS (20 dB under the working level). Each style gives back what it takes.

FIGURE 14 What AMOUNT does: it lowers the compressor's threshold against the working level AUTO LEVEL sets. The six styles differ in ratio, timing and knee; each one's make-up gain gives back what it takes, so AMOUNT changes the hold on the voice, not its loudness.

AMOUNT **is the threshold.** Because AUTO LEVEL brings every voice to the same working level (−18 dBFS RMS), DYNAMICS can be set with one knob. At **0 %** the threshold is −14 dBFS, just over the working level: only the peaks of the loudest words are touched. At **100 %** it is −38 dBFS, 20 dB under the working level: every word is held. Each style has its own **make-up gain**, set from what that style takes from a sung voice at each AMOUNT, so the loudness doesn't change as you turn: within 0.7 dB at 0, 50 and 100 % on a sung voice. A **crest guard** after the make-up holds sudden peaks to 14 dB over the voice's level, so the slower styles' onsets can't become overs.

The six styles:

- **Gentle** (2:1, 20 ms attack, 250 ms release, soft knee): a light, slow compressor that smooths only the loudest words. The default, and the one in most of the effect presets at 45 %.
- **Levelling** (3:1, 30 / 500 ms): an even, steady level, like an optical compressor: the quiet words come up and the loud ones settle. For ballads, spoken voice and anything that should sound untouched.
- **Solid** (4:1, 5 / 120 ms): a firm hold: the voice sits in one place in a dense mix. Rap leads, dance vocals.
- **Shine** (3:1, 10 / 150 ms, plus a lift above 8 kHz that grows with AMOUNT, up to +6 dB): compression with air on top: polished and bright. Modern pop.
- **Punchy** (4:1, 35 / 70 ms): the slow attack lets the start of each word through, the fast release lets go before the next: punch and presence. Rock, soul.
- **Overcooked** (20:1, 0.5 / 50 ms, driven, blended 70 % with the uncompressed voice): squashed and saturated on purpose. Hype ad-libs and aggressive styles.

GATE closes the module down on the quiet parts between phrases: breaths, room, headphone spill. At 0 % it never closes. From there up, its threshold rises from far under any recording's noise to **−30 dBFS at 100 %** (12 dB under the working level); it opens within a millisecond, holds 40 ms, and closes by 30 dB over 150 ms. Set it so the words open it and the gaps don't: 30 to 40 % for a podcast voice with some room (*Podcast Host*, *Clean Interview*). It works after AUTO LEVEL, so the same setting fits a quiet and a loud take.

The GR readout beside GATE shows the compressor's gain reduction right now, in dB (10.0 is 10 dB of reduction). It appears once the reduction is over 0.3 dB. Aim for 3 to 6 dB on the louder words of a lead, 6 to 10 on a rap or broadcast voice.

CHARACTER



FIGURE 15 The CHARACTER and FILTER cards (*Tuned & Robotic: Robot Voice*): 1 AMOUNT with 2 TONE and 3 *Bit Crush* chosen; 4 CUTOFF with 5 RESO, 6 ENV and 7 *Comb* chosen.

Control	Range	Default	What it does
Power	on, off	off	The module on or off.
AMOUNT	0 to 100 % in 1 % steps	40 %	How much colour or distortion.
TONE	0 to 100 % in 1 % steps	50 %	Darker to the left, brighter to the right; 50 % is neutral.
Style	Fuzz, Saturate, Radio, Vinyl, Lo-Fi, Bit Crush	Saturate	Which colour.

CHARACTER colours the voice, from a warm thickening to a broken fuzz. The curves run at **four times the sample rate**, so even the hard styles don't alias into hiss. A slow level match (about a second, up to ±12 dB) keeps the loudness where it was, so AMOUNT changes the colour, not the level; at 0 % the voice passes untouched.

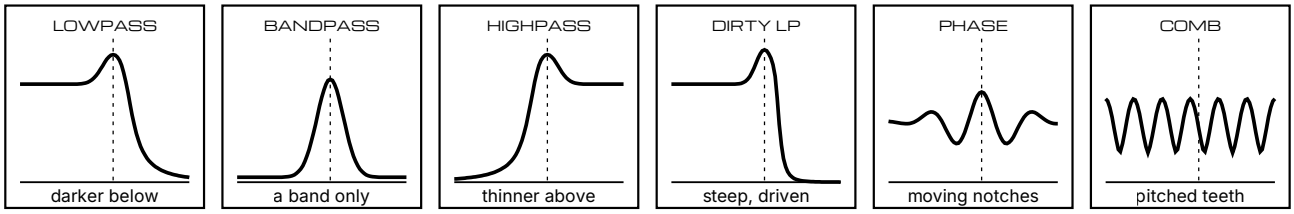
- **Fuzz**: hard, lopsided clipping after a pre-emphasis: buzzing and broken, for screams and breakdowns (*Fuzz Scream*).
- **Saturate**: a warm, round drive: thicker and louder-sounding, never harsh. The default, and the colour in the lead presets (*Warm Lead* 25 %, *Rock Edge* 40 %, *Soul Belter* 35 %).

- **Radio**: band-limited to 400 Hz to 3.5 kHz and driven: a small speaker, a phone line, a megaphone (*Telephone*, *AM Radio*, *Ad-Lib Phone*, *Phone Guest*).
- **Vinyl**: a darker top (rolled off from 18 kHz down towards 6 kHz as you turn), a slow wow of up to 12 cents, rare clicks and a quiet surface hiss: an old record. The crackle is part of the colour and grows with AMOUNT (*Dusty Vinyl*).
- **Lo-Fi**: a lower sample rate (down to a quarter) and fewer bits (down to 8), with the top softened: a dusty, sampled voice (*Bedroom Pop*, *Lo-Fi Tape*).
- **Bit Crush**: far fewer bits (down to 3) and samples (down to a sixteenth of the rate): crunchy digital grit (*8-Bit Crush*, *Robot Voice*).

TONE tilts the colour around 1 kHz, up to 6 dB darker to the left and 6 dB brighter to the right; 50 % is flat. Most of the warm presets sit it at 35 to 45 %; the phone and radio voices at 50 to 60 %.

FILTER

Control	Range	Default	What it does
Power	on, off	off	The module on or off.
CUTOFF	40 Hz to 18 kHz, swept in octaves	2.88 kHz	Where the filter works.
RESO	0 to 100 % in 1 % steps	30 %	A peak at the cutoff: from smooth to whistling.
ENV	−100 to +100 % in 1 % steps	0 %	The voice's level moves the cutoff, up (right) or down (left), by up to 4 octaves.
Style	Lowpass, Bandpass, Highpass, Dirty LP, Phase, Comb	Lowpass	Which filter.



What each style lets through, low frequencies at the left and high at the right; the dotted line is CUTOFF. RESO makes the bump at CUTOFF taller (or PHASE's and COMB's teeth deeper). PHASE sweeps its notches slowly by itself; COMB's teeth sit on a pitch.

FIGURE 16 The six filter shapes, drawn as what each lets through. The dotted line is CUTOFF.

- **Lowpass**, **Bandpass** and **Highpass** are clean 12 dB-per-octave filters whose resonance goes from gentle (RESO 0 %) to whistling (100 %). A low-pass gets darker as CUTOFF goes down; a high-pass thinner as it goes up (*Thin Intro Filter*); a band-pass keeps only the band around CUTOFF (*Telephone Band*).
- **Dirty LP** is a driven four-pole ladder low-pass, 24 dB per octave, whose resonance growls on the low notes (*Dirty Lowpass Growl*).
- **Phase** is a six-stage phaser centred on CUTOFF, swept slowly by itself (one sweep every three seconds or so); RESO is its feedback, which deepens the notches (*Phaser Swirl*).

- **Comb** tunes a comb filter to CUTOFF (40 Hz to 4 kHz): a metallic, pitched resonance, like singing into a pipe; RESO sets how long it rings. Move CUTOFF to tune the metal (*Metal Comb*, *Robot Voice*).

ENV makes the filter follow the voice. To the right (+), the cutoff opens as the voice gets louder: with Lowpass, a wah that opens on every word (*Envelope Wah*, +60 %). To the left (–), it closes on the loud parts instead. ±100 % is a four-octave move.

CUTOFF **and automation.** The big knob is the one to automate: a low-pass opening over a pre-chorus, a high-pass closing on an intro. It is swept in octaves, so the move sounds even from the bottom to the top.

DELAY



FIGURE 17 The DELAY and REVERB cards (*Pop: Radio Pop Lead*): 1 DELAY's MIX, 2 TIME, 3 FEEDBACK and 4 its styles, with *Ping-Pong* chosen; 5 REVERB's MIX, 6 TIME and 7 its styles, with *Plate* chosen.

Control	Range	Default	What it does
Power	on, off	off	The module on or off.
MIX	0 to 100 % in 1 % steps	25 %	How loud the echoes are.
TIME	1/32, 1/16T, 1/16, 1/16D, 1/8T, 1/8, 1/8D, 1/4T, 1/4, 1/4D, 1/2, 1/1	1/8	The echo, in notes of the song's tempo (T triplet, D dotted).
FEEDBACK	0 to 95 % in 1 % steps	35 %	How many times the echo repeats.
Style	Mono, Stereo, Ping-Pong, Grain, Detuned, Filtered	Stereo	Which delay.

The DELAY follows **your DAW's tempo**: TIME is a note value, from a thirty-second to a whole note, with dotted and triplet values in between. A dotted eighth (1/8D) is the classic rolling vocal echo (*Bright Modern Lead*, *Radio Pop Lead*, *Dotted Echo*); a quarter (1/4) the big throw (*Quarter Note Throw*). When the tempo changes the delay glides to the new time rather than jumping. Every repeat passes through a gentle band (150 Hz to 8 kHz), so the echoes sit behind the voice rather than fighting it.

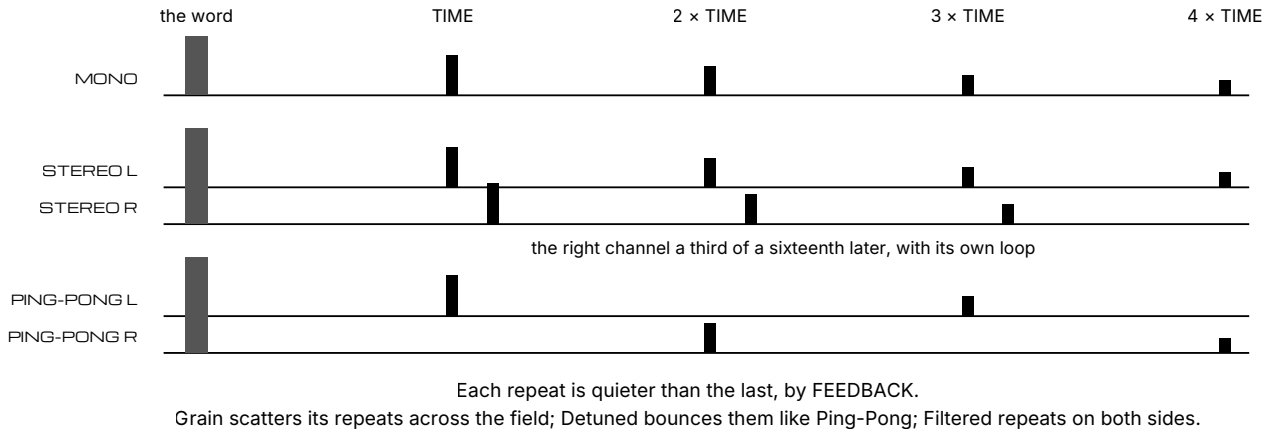


FIGURE 18 Where the echoes land for the first three styles: Mono in the middle; Stereo with the right side a little later than the left; Ping-Pong bouncing from side to side.

- **Mono**: one line, in the middle. The old slap and the tucked echo (*Ballad Hall Lead* 1/4 at 8 %, *Dotted Echo*).
- **Stereo**: the left channel at TIME, the right a third of a sixteenth later, each with its own feedback loop: wide without bouncing. The default.
- **Ping-Pong**: the repeats alternate left and right (*Ping-Pong Eighths*, *Ad-Lib Phone*).
- **Grain**: each repeat is played back as 40 ms grains, scattered up to 30 ms in time, a little in pitch (±0.4 semitone) and across the stereo field: a shimmering, textured trail (*Grain Cloud*).
- **Detuned**: the repeats bounce from side to side like Ping-Pong, and each pass is detuned by 12 cents (up on the way left, down on the way right), so they shimmer and widen as they go (*Detuned Echo*, *Late Night Falsetto*).
- **Filtered**: each repeat passes a resonant band-pass that sweeps slowly by itself, from 500 Hz to 2.8 kHz and back about every seven seconds: the dub echo that gets darker and thinner as it repeats (*Filtered Dub Echo*).

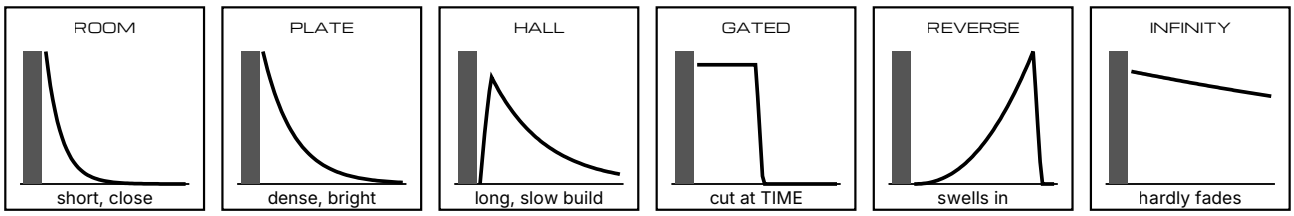
A change of style fades the old echoes out, clears them and fades the new style in over 20 ms, so you can audition the six while the song plays. FEEDBACK sets how many repeats you hear: 20 to 30 % gives two or three, 60 % a long trail, 95 % a trail that goes on for many bars. MIX is the echoes' level against the voice; the voice itself stays at full level up to 100 % (where it is 4.6 dB down, never gone), so the delay can't take the voice away.

> **Tip:** For a throw on one word, automate DELAY's MIX from 0 % up on that word and back down after it: the repeats already in the air keep going. On a mono track, *Stereo*, *Ping-Pong* and *Detuned* fold into the middle and sound like *Mono*.

REVERB

Control	Range	Default	What it does
Power	on, off	on	The module on or off.
MIX	0 to 100 % in 1 % steps	25 %	How much space around the voice.
TIME	0.3 to 8 s	1.1 s	How long the reverb lasts.
Style	Room, Plate, Hall, Gated, Reverse, Infinity	Plate	Which space.

The REVERB is **VEIRA Space's engine**, set up for a voice: a 15 ms pre-delay so the words stay clear, the reverb's low end cut at 150 Hz and its top at 12 kHz so it never clouds the voice, and six spaces.



The level of the reverb after a word (the grey block), over time. TIME sets how long ROOM, PLATE and HALL ring, how long GATED stays open before it is cut, and how long REVERSE takes to swell. INFINITY ignores TIME: it lasts about 20 s.

FIGURE 19 How each space decays after a word.

- **Room**: a small, close room; TIME runs up to 2.5 s here, since a longer room isn't one (*Small Room*, *Intimate Close Lead*).
- **Plate**: a bright, smooth plate: the classic vocal reverb and the default. 0.8 to 1.3 s on a lead, 1.5 to 1.8 s for a bigger chorus (*Vintage Plate*).
- **Hall**: a large hall with a slow build (*Ballad Hall Lead*, *Big Hall*, *Gospel Lead*).
- **Gated**: a big, dense reverb cut short; TIME sets how long it stays open, from 0.15 to 0.75 s: the eighties (*Gated Eighties*).
- **Reverse**: the reverb played backwards, in slices as long as TIME sets (a quarter of a second to one and a half), so it swells up into the end of each phrase (*Reverse Swell*).
- **Infinity**: a cathedral that hardly decays (about 20 s), with a little shimmer an octave up: an endless pad under the voice. It ignores TIME (*Infinite Wash*).

MIX blends the reverb with the voice with equal power: at 100 % the voice is 3 dB down and the reverb full, so even an all-wet setting keeps the words. Leads sit at 10 to 20 %, backing parts at 20 to 35 %, the *Spaces* presets at 30 to 45 %.

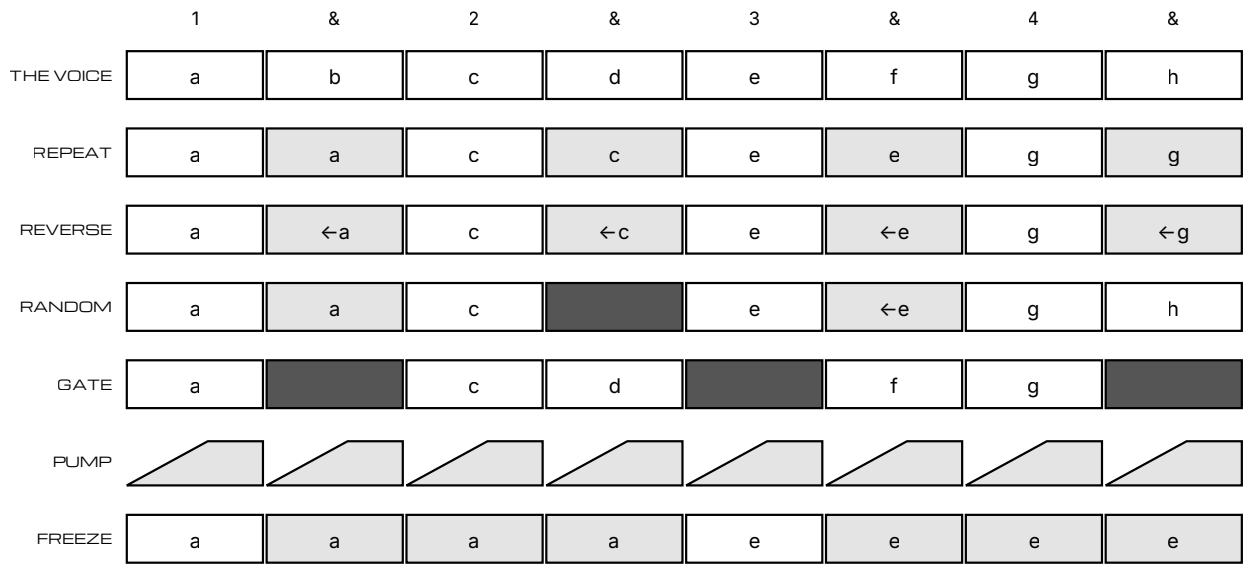
CHOP



FIGURE 20 The CHOP card (*Chop & Glitch: Reverse Chops*): 1 AMOUNT, 2 PRE (on), 3 RATE and 4 the styles, with *Reverse* chosen. 5 The thread between the REVERB and CHOP cards lights while PRE is on, to show that CHOP now runs before the DELAY and REVERB.

Control	Range	Default	What it does
Power	on, off	off	The module on or off.
AMOUNT	0 to 100 % in 1 % steps	50 %	How much of the chopped voice against the live one.
RATE	1/32 to 1/1, as DELAY's TIME	1/16	The slice length, in notes of the song's tempo.
PRE	on, off	off	CHOP before the DELAY and REVERB (their tails flow on); off, after them (the tails are chopped too).
Style	Repeat, Reverse, Random, Gate, Pump, Freeze	Gate	Which rhythm.

CHOP cuts the voice into slices of RATE's length and does something to every slice, **in time with the song**. While your DAW plays, the slices are locked to its bars and beats; when it is stopped, CHOP runs on at the song's tempo so you can audition it. The slices are placed so that they land exactly on the beat once your DAW has made up the plugin's latency, and every slice edge is crossfaded over 3 ms, so the cuts are clean but not clicky.



One bar at RATE 1/8. White slices are the live voice, grey ones are replayed copies, dark ones silence; ← plays the copy backwards.

FIGURE 21 What each style does to a bar of the voice, cut into eighth-note slices a to h.

- **Repeat**: every other slice replays the slice before it: the classic stutter (*Stutter Repeat*, 1/16 at 100 %).
- **Reverse**: every other slice plays the slice before it backwards (*Reverse Chops*, 1/8).
- **Random**: each slice, by chance, plays as it is (40 %), repeats the one before (25 %), plays it backwards (20 %) or drops out (15 %): a glitch that never repeats itself (*Random Glitch*).
- **Gate**: a sixteen-step pattern opens and closes the voice: the trance gate (*Sixteenth Gate*, 1/16). The pattern is fixed; RATE sets how fast it runs.
- **Pump**: the voice ducks at the start of each slice and swells back over the first 60 % of it, like a pad under a sidechain. RATE 1/4 pumps on every beat (*Sidechain Pump*).
- **Freeze**: every fourth slice is held and looped under the next three: a drone from a moment of the voice (*Freeze Drone*, 1/4).

AMOUNT blends the chopped voice with the live one. At 100 % you hear only the chop; at 50 % the live voice carries on underneath and the chop becomes a texture. RATE sets the slice: 1/16 for stutters and gates, 1/8 for reverses, 1/4 for pumps and freezes.

PRE moves CHOP in front of the DELAY and REVERB. Off (the default), CHOP comes after them and chops their tails too: the whole sound stutters, which is what you want for a gate or a pump. On, the echoes and the reverb are fed the chopped voice and ring on over it: a reversed slice leaves a plate behind it (*Reverse Chops*), a frozen slice becomes a drone in a hall (*Freeze Drone*). A dashed thread between the REVERB and CHOP cards lights while CHOP is on with PRE. Changing it fades the output out and back in over 15 ms, so you can switch it while the song plays.

Tip: Automate the CHOP power switch, or AMOUNT, to bring a stutter in on the last beat of a phrase and take it out again. The slices stay on the grid wherever you switch it on.

The global bar



FIGURE 22 The left of the global bar: 1 INPUT, 2 AUTO LEVEL with its gain readout (here +0.7 dB), 3 LOW CUT and HIGH CUT, and 4 the start of the MIX slider.



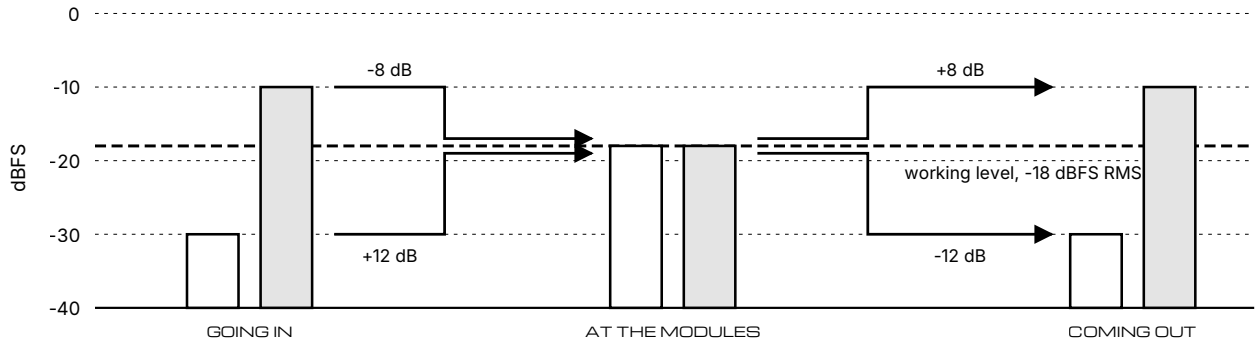
FIGURE 23 The right of the global bar: 4 the end of the MIX slider (here at 100 %), 5 LOCK, 6 DOUBLING, FOCUS EQ and GLUE, and 7 OUTPUT.

Control	Range	Default	What it does
INPUT	−24 to +24 dB in 0.1 dB steps	0.0 dB	The level going in.
AUTO LEVEL	on, off	on	Brings the voice to the working level before the modules and gives the gain back at the end.
LOW CUT	Off (20 Hz) to 600 Hz	Off	Removes rumble under this frequency.
HIGH CUT	2 kHz to Off (22 kHz)	Off	Removes hiss above this frequency.
MIX	0 to 100 % in 1 % steps	100 %	The whole effect against the dry voice, in time with it.
LOCK	on, off	off	Presets leave MIX where it is. Not automatable, not saved in presets.
DOUBLING	0 to 100 % in 1 % steps	0 %	A second take made from the first, left and right.
FOCUS EQ	0 to 100 % in 1 % steps	0 %	Brings the voice forward: presence and air, less mud.
GLUE	0 to 100 % in 1 % steps	0 %	A gentle compressor on the whole result.
OUTPUT	−24 to +12 dB in 0.1 dB steps	0.0 dB	The level going out.

INPUT, LOW CUT and HIGH CUT

INPUT sets the level going into everything, including the dry copy that MIX uses. With AUTO LEVEL on you rarely need it. LOW CUT takes out rumble, pops and the proximity boom under the voice: 80 to 120 Hz on a lead, 150 to 300 Hz on backing parts you want to tuck, 350 Hz for a telephone. HIGH CUT takes hiss and harshness off the top: most leads leave it off; 12 to 16 kHz softens a bright voice, 3.5 to 5 kHz makes a phone or a radio. Both are gentle 12 dB-per-octave filters; turned all the way down or up they're off, and the readout says so.

AUTO LEVEL: the same treatment for every take



A quiet take (white, -30 dBFS) and a loud one (grey, -10 dBFS) both reach the modules at -18 dBFS, and each comes out at its own level.

FIGURE 24 AUTO LEVEL: the gain goes in before the modules and comes back out after them.

Every compressor, gate, saturator and filter reacts to the level it is fed. AUTO LEVEL takes that variable away: it measures the voice (RMS over 0.4 s, ignoring what's under 100 Hz) and moves a slow gain, within ±18 dB, until the voice sits at about **−18 dBFS RMS**. It moves up over about 0.3 s and down over about a second, and it holds still while the voice is quiet (under −50 dBFS), so it doesn't pump up the noise between phrases. **The same gain is taken back at the end**, delayed to match the voice, so the output follows the input: a quiet take comes out quiet and a loud one loud, each having had the same treatment. The readout under the switch shows the gain in use right now (+4.7 dB on the voice the presets were measured on). It is on in every factory preset, and in Init.

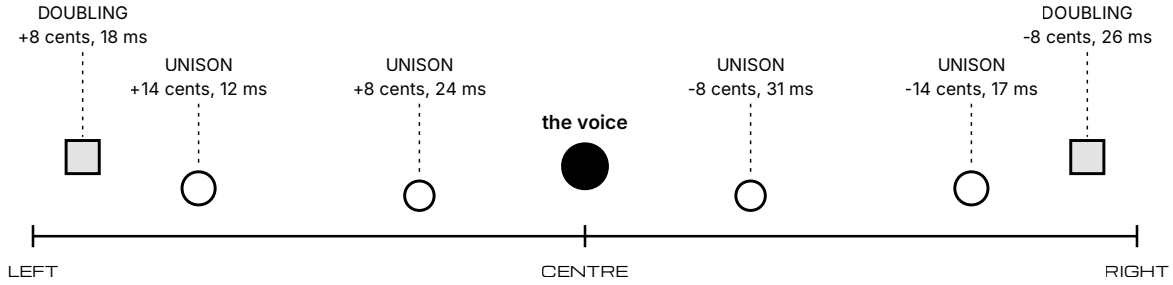
Switch it **off** when you want the modules to react to the real level, for example to drive DYNAMICS harder on the loud lines on purpose. Then INPUT is your gain stage: the presets were made for a voice at about −18 dBFS RMS, so set INPUT for that, or expect them to do less on a quiet take and more on a loud one.

MIX and LOCK

MIX blends the whole effect with the dry voice. The dry copy is delayed by exactly the plugin's latency, so the two stay in time and add up cleanly: at 0 % the output is the input, bit for bit. Use it to blend a heavy effect under the natural voice (a crushed or chopped voice at 40 %), or to pull a preset back a little as a whole. The caption under the slider shows its value while you hover or drag.

LOCK keeps MIX where it is while you change presets. Normally a preset sets everything, MIX included (every factory preset has it at 100 %). With LOCK on, presets leave MIX alone, so you can set a blend of, say, 50 % and then step through the factory presets at that blend. LOCK is a convenience, not a sound: it isn't automatable and isn't saved in presets, though it is saved with your session.

DOUBLING, FOCUS EQ and GLUE



UNISON's four copies at 100 % (circles) and DOUBLING's two takes (squares), each a little later than the voice and a little out of tune with it.

FIGURE 25 The stereo field: the voice in the middle, UNISON's copies around it and DOUBLING's takes hard left and right. On a mono track they all fold into the middle.

DOUBLING makes a second take from the first: two copies, one 18 ms behind the voice and 8 cents sharp, panned left, the other 26 ms behind and 8 cents flat, panned right. It is the sound of a singer doubling their own line, without the second take. 15 to 35 % thickens a lead a little (*Pop Sheen* 25 %, *Wide Chorus Vocal* 35 %); 40 to 60 % is a wide double (*Tight Double* 60 %, *Gang Vocals*). Unlike UNISON, which sits around the voice and is made before the modules, DOUBLING is added at the very end, after the effects, so it doubles the echoes and the room too.

FOCUS EQ is a presence curve in one knob: up to **+5 dB at 3.2 kHz**, **−3 dB at 320 Hz** and **+2.5 dB of air above 10 kHz**, all scaled by the knob. It brings a voice forward in a busy mix without a separate EQ. 15 to 25 % is a nudge, 35 to 50 % a clearly forward pop or broadcast voice (*Dance Pop Vocal* 50 %, *Radio Announcer* 40 %).

GLUE is a gentle 2:1 compressor (10 ms attack, 120 ms release, soft knee) on the whole result, voice and effects together, so they sit as one thing. The knob lowers its threshold from −6 to −28 dB; its level is matched to what went in, so it changes the hold, not the loudness. 20 to 35 % on a lead that has delays and reverb, 40 % on a dense broadcast voice.

OUTPUT sets the level leaving the plugin, after MIX. The factory presets use it to land each preset at the input's loudness; on your own voice, trim it until the OUT meter matches IN with the plugin bypassed and switched in.

Displays

VEIRA Vocal Studio has no separate display window; what it is doing is shown on the controls themselves.

- **The NOTE readout** (in the VOICE section) shows the note being sung, and with CORRECTION above 0 % an arrow and the note it goes to. -- means no pitched note is being heard. See *Controls ▸ VOICE ▸ CORRECTION* and its figure.
- **The GR readout** (on the DYNAMICS card, beside GATE) shows the compressor's gain reduction in dB. It appears once the reduction is over 0.3 dB. See *Controls ▸ DYNAMICS* and its figure.
- **The AUTO LEVEL readout** (under its switch) shows the gain AUTO LEVEL is applying right now, in dB. A big positive number means a quiet take; a negative one a hot take.
- **The orbs** glow with the voice's level while VOICE is on; a dark orb means no signal, or VOICE off.
- **Every knob's value** appears in place of its label while you hover over it or drag it, with its unit (dB, st, %, Hz or kHz, ms or s, or a note value). The MIX slider's value appears under it the same way. To set an exact value, right-click and type it.
- **The arcs** around the knobs and orbs show each setting at a glance, lit from the start of the travel (or from the top, for PITCH and FORMANT).

The IN and OUT meters

The slim meters on either side of the module cards show the level going into the plugin (IN, after INPUT) and coming out of it (OUT, after OUTPUT), from −60 to +6 dBFS, one bar per channel. The brighter body of each bar is the RMS level, the fainter extension above it the peak (falling at 26 dB per second), and the thin line at the top the highest recent peak, held for a second and a half. The line turns red (on screen) when a peak reaches 0 dBFS. Hover over a meter for the held peak as a number. On a mono track both bars show the same signal.

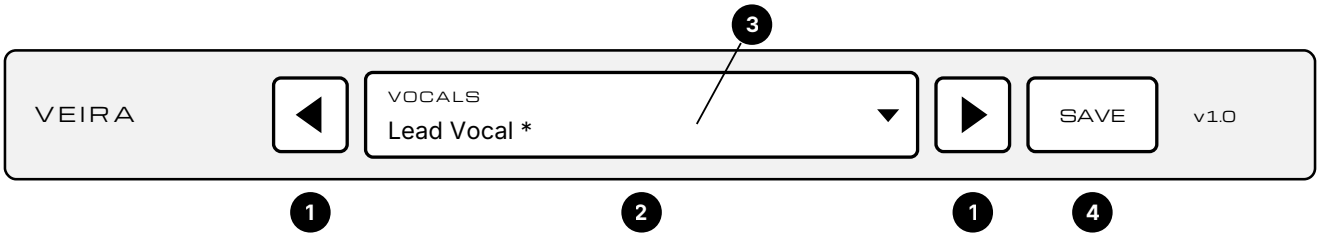
Because AUTO LEVEL gives back its gain and the modules and presets are level-matched, OUT normally reads close to IN; if it reads much higher, look at OUTPUT and at DELAY's or REVERB's MIX.

Presets

VEIRA Vocal Studio comes with **80 factory presets** in 13 folders. Each one is a starting point made for a real job. Load one close to what you need, then adjust the first control its note mentions.

The preset browser

The preset browser sits in the header, at the top of the window.



The header: logo and plugin name, then the preset browser, then the version.

FIGURE 26 The preset browser.

- 1 ◀ ▶ step to the previous or next preset, through every folder in order.
- 2 **The display** shows the folder (small, above) and the preset's name. Click it to open the preset menu. Hover over it to read the preset's note: what it's for and what to adjust first.
- 3 **An asterisk (*)** after the name means you've changed something since the preset loaded.
- 4 SAVE stores your current settings as a user preset.

The preset menu

Clicking the display opens a menu with:

- one submenu per **factory folder**, in the order shown below;
- **User Presets**, with your own presets in their folders;
- **Save...**, **Delete** (only for user presets: factory presets can't be deleted), **Rescan** (re-reads the user folder, if you copied presets in by hand) and **Show User Preset Folder** (opens it in Finder or Explorer).

Saving your own presets

1. Set VEIRA Vocal Studio the way you like it.
2. Press SAVE. A dialog opens inside the plugin window.
3. Type a name, choose a folder or type a new folder name, and confirm.

Each user preset is one small .veirapreset file:

	Folder
macOS	~/Library/Audio/Presets/VEIRA/VEIRA Vocal Studio/<Folder>/
Windows	%APPDATA%\VEIRA\VEIRA Vocal Studio\Presets\<Folder>\

Copy these files to move your presets to another computer, or to back them up. Presets keep loading in future versions: any control added later simply starts at its default.

Presets in your DAW

The factory presets also appear in your DAW's own preset or program menu, named *Folder: Name* (for example *Vocals: Lead Vocal*). Whatever preset is selected is saved with your session and comes back when you reopen it, including the asterisk if you'd changed it.

The **Default ▸ Init** preset equals a fresh instance of the plugin: every control at its default.

The factory folders

Folder	Presets
Default	Init
Lead Vocals	Polished Lead, Warm Lead, Bright Modern Lead, Upfront Dry Lead, Ballad Hall Lead, Intimate Close Lead, Rock Edge
Pop	Radio Pop Lead, Pop Sheen, Dance Pop Vocal, Bedroom Pop, Tuned Pop Lead, Wide Chorus Vocal
Hip-Hop & Rap	Rap Lead Upfront, Trap Vocal, Boom Bap Verse, Ad-Lib Phone, Dark Drill Vocal, Hype Ad-Libs
R&B & Soul	Silky R&B Lead, Neo-Soul Warmth, Modern R&B Gloss, Soul Belter, Late Night Falsetto, Gospel Lead
Backing & Doubles	Wide Stack, Tight Double, Airy Background Pad, Octave Stack, Gang Vocals, Tucked Harmonies
Tuned & Robotic	Hard Tune Major, Hard Tune Minor Pent, Natural Tune, Glassy Pop Tune, Robot Voice, Synth Voice
Pitch & Formant	Octave Down Double, Octave Up Sparkle, Fifth Harmony, Monster Voice, Chipmunk, Slowed Down, Small Bright Voice
Lo-Fi & Character	Telephone, AM Radio, Dusty Vinyl, Lo-Fi Tape, 8-Bit Crush, Fuzz Scream
Filter FX	Envelope Wah, Telephone Band, Thin Intro Filter, Dirty Lowpass Growl, Phaser Swirl, Metal Comb
Delay Throws	Quarter Note Throw, Ping-Pong Eighths, Dotted Echo, Grain Cloud, Detuned Echo, Filtered Dub Echo
Spaces	Small Room, Vintage Plate, Big Hall, Gated Eighties, Reverse Swell, Infinite Wash
Chop & Glitch	Sixteenth Gate, Sidechain Pump, Stutter Repeat, Reverse Chops, Random Glitch, Freeze Drone
Podcast & Voice	Podcast Host, Warm Narrator, Clean Interview, Radio Announcer, Phone Guest

The folders are sorted by job:

- **The lead folders** (*Lead Vocals*, *Pop*, *Hip-Hop & Rap*, *R&B & Soul*) are finished vocal sounds: DYNAMICS at 35 to 60 %, a plate or a hall at 8 to 20 %, often a tempo-synced echo at 8 to 14 %, and FOCUS EQ and GLUE where the style wants them. They compress 3 to 8 dB on average (*Hype Ad-Libs* more, on purpose), and each comes out at the loudness it went in.
- **Backing & Doubles** are parts behind a lead: UNISON, DOUBLING and a higher LOW CUT.

- **Tuned & Robotic** and **Pitch & Formant** are the VOICE section's folders: natural and hard correction, harmonies, octaves and character voices. The tuned presets use **example keys**: set KEY and SCALE to the song's first.
- **Lo-Fi & Character**, **Filter FX**, **Delay Throws**, **Spaces** and **Chop & Glitch** each put one module forward, with Gentle DYNAMICS at 45 % and a modest plate behind it. Their notes say which control to move first; the throws and the chops work on the song's tempo.
- **Podcast & Voice** is spoken voice: dry, with GATE where a room needs it.

Each preset's note (hover over the preset's name) says what it's for and which control to move first.

Levels. Each preset was measured on the real plugin on a sung vocal that suits its folder (a female voice for most, a male voice for the hip-hop and podcast folders, a harsh, bright one for the two rock presets that say so), peaking at −6 dBFS and at 120 BPM. OUTPUT is set so each comes out within 0.5 dB of the input's loudness, and no peak goes over −1 dBFS, so you can compare with bypass fairly. One is quieter on purpose and its note says so: *Telephone Band*, about 3 dB under, because a thin band reads louder than it measures. Because AUTO LEVEL is on in every preset, they behave the same on a quiet take: on the same voice 12 dB down, the lead presets come out within about 1 dB of it and compress the same amount.

Recipes

Step-by-step starting points. In each, load the preset, then compare with bypass: the OUT meter should read about what IN does.

A finished lead vocal

1. Insert VEIRA Vocal Studio on the vocal track (a mono-to-stereo track if the recording is mono) and load a lead from your style's folder, for example **Lead Vocals ▸ Polished Lead** or **Pop ▸ Radio Pop Lead**.
2. Watch the GR readout: set DYNAMICS' AMOUNT for 4 to 6 dB of reduction on the louder words.
3. Try the other DYNAMICS styles while it plays: *Levelling* for a smoother ballad, *Shine* for a brighter pop lead, *Solid* for a dense mix.
4. Set REVERB's MIX for the room and, if the preset has one, DELAY's MIX for the echo. Turn FOCUS EQ up a little if the voice sits behind the guitars.
5. If the whole thing is a touch much, pull MIX back to 80 or 90 % rather than lowering every module.

Natural pitch correction

1. Load **Tuned & Robotic ▸ Natural Tune** (CORRECTION 30 %).
2. Set KEY and SCALE to the song's. If the melody uses notes outside the key, choose *Chromatic*.
3. Play the vocal and watch the NOTE readout: the right-hand note should be the one the melody wants, and move with it.
4. If you can hear the correction, lower CORRECTION towards 20 %; if held notes still drift, raise it towards 45 %.

The hard tune effect

1. Load **Tuned & Robotic ▸ Hard Tune Major**, or **Hard Tune Minor Pent** for melodic rap (five safe notes).
2. Set KEY and SCALE to the song's, or every note jumps to the wrong place.
3. Leave CORRECTION at 100 %. For more of the effect, add UNISON 20 % (*Glassy Pop Tune*) or a dotted echo.
4. For the robot, load **Robot Voice** instead and move FILTER's CUTOFF to tune the metal.

A harmony or an octave from one voice

1. Duplicate the lead track, or send it to a stereo aux, and insert the plugin there. Load **Pitch & Formant ▸ Fifth Harmony** (PITCH +7 st, VOICE MIX 45 %).
2. Set PITCH to the interval you want: +7 a fifth up, +5 a fourth, +12 an octave, −12 an octave down (*Octave Down Double*), −5 a fourth below.
3. Set VOICE MIX for the harmony's level under the natural voice (or to 100 % if the natural voice is already on the other track).
4. Keep FORMANT at 0 for a natural second singer, or move it ±2 to ±3 st so the harmony sounds like somebody else.
5. Add UNISON 20 to 40 % to widen the harmony.

A character voice

1. Load **Pitch & Formant ▸ Monster Voice**, **Chipmunk**, **Slowed Down** or **Small Bright Voice**.
2. For a creature, move FORMANT first: down is bigger. Then PITCH, a few semitones.
3. For the sped-up or slowed-down tape, keep LINK on and move only PITCH: +5 to +12 for the chipmunk, −3 to −7 for the screwed voice.
4. Add CHARACTER (*Saturate* 20 to 30 %) for grit, and a *Room* at 10 to 15 % to put the creature somewhere.

Backing vocals behind the lead

1. On the backing vocal track, load **Backing & Doubles ▸ Wide Stack** (or **Tucked Harmonies** for parts that should sit well under the lead).
2. Set LOW CUT first (150 to 300 Hz), then HIGH CUT if the stack is too bright, to make room for the lead.
3. Set UNISON for the size of the group and DOUBLING for the width.
4. For a double of the lead itself, use **Tight Double** on a copy of the lead track and bring it in under the lead.

A delay throw on one word

1. Load **Delay Throws ▸ Quarter Note Throw** (or *Dotted Echo* for a rolling one) and set TIME to the note you want.
2. Set FEEDBACK for how long the throw repeats (45 % is three or four repeats).
3. Automate DELAY's MIX: 0 % through the line, up to 30 or 40 % on the last word, back to 0 % after it. The echoes already in the air keep repeating.
4. For a throw that darkens as it goes, choose *Filtered*; for one that shimmers, *Detuned* or *Grain*.

Stutters and chops on the beat

1. Load **Chop & Glitch ▸ Stutter Repeat** and play the song: the stutter locks to your DAW's grid.
2. Set RATE: 1/16 for a tight stutter, 1/8 for a slower one. Try *Reverse* and *Random*.
3. Automate the CHOP power switch, or AMOUNT, so the chop happens only on the fills.
4. For a trance gate or a sidechain pump on a pad-like vocal, use **Sixteenth Gate** or **Sidechain Pump** and leave PRE off so the reverb pumps too; for reversed slices with a tail, **Reverse Chops** has PRE on.

A podcast or voice-over

1. Load **Podcast & Voice ▸ Podcast Host** (or **Clean Interview** for a guest with room noise).
2. Watch the GR readout: set AMOUNT for 5 to 8 dB of reduction on the louder sentences.
3. Set GATE so the pauses close it and the words open it: raise it until the room drops out between sentences, then back off a little.
4. FOCUS EQ 25 to 35 % for clarity; GLUE 30 to 40 % for a dense broadcast voice (*Radio Announcer*).

Specifications

Type	Creative vocal plugin: a VOICE section (pitch, formant, unison, correction) and six style-based modules (DYNAMICS, CHARACTER, FILTER, DELAY, REVERB, CHOP) with a global bar
Channels	Mono, stereo, and mono in to stereo out. UNISON, DOUBLING, the stereo delays and the reverb spread a mono voice
Latency	2,943 samples at 48 kHz (61 ms); 2,706 at 44.1 kHz; 5,855 at 96 kHz; 11,678 at 192 kHz. Reported to the host; constant whichever modules are on
Transparency	Every module off, MIX at 0 %, or bypassed: the output is the input delayed by the latency, bit for bit
Switching	A module on or off crossfades over 20 ms; a change of style switches smoothly; PRE fades the output out and in over 15 ms; knobs glide over 30 ms
INPUT, OUTPUT	−24 to +24 dB and −24 to +12 dB, in 0.1 dB steps (both 0.0 dB)
AUTO LEVEL	On by default. Toward −18 dBFS RMS (0.4 s, ignoring what's under 100 Hz), within ±18 dB, 0.3 s up and 1 s down, held under −50 dBFS; the same gain taken back at the end, in time with the voice
LOW CUT, HIGH CUT	Off (20 Hz) to 600 Hz; 2 kHz to Off (22 kHz); both 12 dB/oct, both Off by default
VOICE	PITCH and FORMANT −12 to +12 st in 0.01 st steps (0), LINK (off), UNISON 0–100 % (0), CORRECTION 0–100 % (0), KEY C to B (C), SCALE Chromatic, Major, Minor, Major Pentatonic, Minor Pentatonic (Chromatic), VOICE MIX 0–100 % (100). A pitch-synchronous shifter on the mid (75 Hz to 1 kHz tracked); the sides pass delayed. Measured: +7 st on 220 Hz gives 328.8 Hz; FORMANT ±6 st moves the timbre's centroid from 775 Hz to 1,080 / 545 Hz; a note 40 cents flat lands within 2 cents at 100 %
CORRECTION	0–25 %: the amount, from 0 to all the way; 25–100 %: the glide, 225 ms at 25 %, 100 ms at 50 %, 25 ms at 75 %, instant at 100 %
UNISON	Up to four copies, 12 / 17 / 24 / 31 ms behind the voice; to 50 % one pair (±4 to ±9 cents, panned wide), to 100 % a second pair nearer the middle, the detuning to ±14 and ±8 cents
DYNAMICS	AMOUNT 0–100 % (50): threshold from −14 to −38 dBFS; GATE 0–100 % (0): off, then a threshold up to −30 dBFS, 1 ms open, 40 ms hold, 30 dB closed over 150 ms; styles Gentle 2:1 20/250 ms, Levelling 3:1 30/500, Solid 4:1 5/120, Shine 3:1 10/150 with up to +6 dB over 8 kHz, Punchy 4:1 35/70, Overcooked 20:1 0.5/50 driven at 70 % wet; make-up per style; a 14 dB crest guard; the GR readout. On by default (Gentle 50 %)
CHARACTER	AMOUNT 0–100 % (40), TONE 0–100 % (50); Fuzz, Saturate, Radio (400 Hz–3.5 kHz), Vinyl, Lo-Fi, Bit Crush; 4x oversampled; level-matched (about 1 s, ±12 dB). Off by default
FILTER	CUTOFF 40 Hz–18 kHz in octaves (2.88 kHz), RESO 0–100 % (30), ENV −100 to +100 % (0), up to 4 octaves; Lowpass, Bandpass, Highpass 12 dB/oct (Q 0.6 to 12), Dirty LP 24 dB/oct ladder, Phase 6 stages swept at 0.3 Hz, Comb 40 Hz–4 kHz. Off by default

DELAY	MIX 0–100 % (25), TIME 1/32 to 1/1 with dotted and triplet values (1/8), FEEDBACK 0–95 % (35); Mono, Stereo, Ping-Pong, Grain, Detuned (±12 cents), Filtered (sweep at 0.15 Hz); repeats band-limited 150 Hz–8 kHz; the voice −4.6 dB at MIX 100 %; up to 4.5 s of delay. Off by default
REVERB	MIX 0–100 % (25), TIME 0.3–8 s (1.1 s); Room (to 2.5 s), Plate, Hall, Gated (open 0.15–0.75 s), Reverse (slices 0.25–1.5 s), Infinity (about 20 s, shimmer); 15 ms pre-delay, 150 Hz–12 kHz; equal-power blend, the voice −3 dB at MIX 100 %. On by default (Plate, 25 %)
CHOP	AMOUNT 0–100 % (50), RATE 1/32 to 1/1 (1/16), PRE (off); Repeat, Reverse, Random, Gate (16 steps), Pump, Freeze; locked to the host's beat, set back by the latency; free-running at the tempo when stopped; 3 ms slice edges. Off by default
DOUBLING	0–100 % (0): two copies, ±8 cents, 18 and 26 ms, left and right
FOCUS EQ	0–100 % (0): up to +5 dB at 3.2 kHz, −3 dB at 320 Hz, +2.5 dB over 10 kHz
GLUE	0–100 % (0): 2:1, 10/120 ms, soft knee, threshold −6 to −28 dB, level-matched
MIX, LOCK	0–100 % (100), the dry voice delayed to stay in time; LOCK off, not automatable, not in presets
Metering	IN and OUT, −60 to +6 dBFS, peak and RMS per channel with peak hold; the GR, AUTO LEVEL and NOTE readouts
Presets	80 factory presets in 13 folders, plus your own
Window	1320 × 860 points; resizable from 75 % to 200 %, saved with the session
Formats	VST3, Audio Units, standalone; AAX once signed
Plugin ID	Manufacturer Veir, plugin Vvsd

Troubleshooting and questions

The sound drops out for a moment every 30 seconds. The plugin isn't activated on this computer. Open VEIRA Manager and sign in; the strip at the bottom of the plugin's window says exactly what's wrong. See *Getting started ▸ Activation and the licence*.

I don't see VEIRA Vocal Studio in my DAW. Make sure VEIRA Manager shows SHOW (installed) for it, then rescan plugins in your DAW. In Logic Pro, open the Plug-in Manager and choose *Reset & Rescan Selection* if needed. Look under the manufacturer **VEIRA**.

The singer hears themselves late while recording. VEIRA Vocal Studio has a latency of about 61 ms, which is too much to sing through. Bypass it (or switch the track's plugins off) while you record, and switch it back on for playback. Your DAW keeps the recorded take in time.

The voice is out of time with the other tracks. Your DAW compensates for the plugin's latency on playback; check that its plugin delay compensation is on. The latency is the same whichever modules are on, so switching modules never shifts the voice.

CORRECTION pulls notes to the wrong ones. Set KEY and SCALE to the song's key. The factory presets use example keys. If the melody uses notes outside the scale, use *Chromatic*. Watch the NOTE readout: the right-hand note is where each note goes.

CORRECTION sounds robotic when I want it natural. Lower it: 25 to 40 % is natural, 100 % is the effect on purpose. Below 25 % it also takes less of the way, so a little drift stays.

UNISON and DOUBLING don't get wider. On a mono track the left and right are summed into one channel, so the copies fold into the middle and only thicken. Insert the plugin on a mono-to-stereo track or a stereo aux, and the voice spreads.

The chops don't land on the beat. CHOP locks to your DAW's bars while it plays; make sure the transport is running (when it's stopped, CHOP free-runs at the tempo). Check that your DAW's plugin delay compensation is on: the slices are placed to land on the beat after the latency is made up.

The echoes aren't in time with the song. TIME is a note value at your DAW's tempo, so check the song's tempo and whether the track has a tempo map. There is no millisecond mode; for a free time, use a dotted or triplet value.

A preset doesn't change MIX. LOCK is on. Switch it off and presets set MIX again (every factory preset has it at 100 %).

The presets do too much, or too little, on my recording. With AUTO LEVEL on they shouldn't: it brings any take to the working level. If you switched it off, set INPUT so the voice reads about −18 dBFS RMS on the IN meter, or switch it back on.

I hear a small change in the attack of words with the VOICE section on. Like every pitch shifter of this kind, the shifter works one period of the voice at a time and can move a note's start by up to half a period, about 4 ms on a low voice. Switch VOICE off where it isn't needed, or automate it.

Breaths and unpitched sounds aren't shifted. The shifter follows pitched notes (75 Hz to 1 kHz); breaths, "s" sounds and noise have no note and pass through with the voice's timing. That is normal and keeps consonants crisp.

The level jumps when I switch a module on. It shouldn't: every module is level-matched and switching crossfades. If it does, look at the module's MIX (DELAY and REVERB add energy at high settings), at GATE (it may be closing on a quiet take) and at OUTPUT.

The window is too small (or too big) on my screen. Drag its bottom-right corner. It scales from 75 % to 200 % and keeps its proportions; the size is saved with your session.

How do I type an exact value? Right-click the knob or orb (Control-click on a Mac), type the number and press Return.

Do my sessions keep my settings? Yes. Every control, the selected preset and the window size are saved with your session. Every control can be automated except LOCK.

Contact: support@veiraaudio.com. Mention your computer (Mac or Windows and its version), your DAW and version, and what you see.



VEIRA Vocal Studio · User Manual · Version 0.2.0

Support: **support@veiraaudio.com** · Web: **veiraaudio.com**

© 2026 VEIRA. Proudly made in the Dominican Republic. VST is a trademark of Steinberg Media Technologies GmbH. Audio Units, Logic Pro and macOS are trademarks of Apple Inc. Pro Tools and AAX are trademarks of Avid Technology, Inc. All other trademarks belong to their owners and are used only to describe compatibility.