

Modern C++ Cheat Sheet

1. Common Data Types

Type	#include	Description
<code>int</code>	built-in	whole numbers
<code>double</code>	built-in	floating-point numbers
<code>float</code>	built-in	smaller floating-point (suffix <code>f</code>)
<code>char</code>	built-in	single character (numeric type, 'A' = 65)
<code>bool</code>	built-in	<code>true</code> / <code>false</code>
<code>long</code>	built-in	at least as large as <code>int</code> (suffix <code>l</code>)
<code>long long</code>	built-in	larger than <code>long</code> (suffix <code>ll</code>)
<code>unsigned int</code>	built-in	non-negative whole numbers (suffix <code>u</code>)
<code>size_t</code>	<code><cstdint></code>	unsigned type for sizes, returned by <code>.size()</code>
<code>std::string</code>	<code><string></code>	dynamic character string
<code>std::string_view</code>	<code><string_view></code>	non-owning read-only view of a string
<code>auto</code>	keyword	compiler deduces the type

Traps: `auto x = "hello"` gives `const char*`, not `std::string` – use `"hello"s` suffix with `using namespace std::string_literals`. Brace initialization `int x{3.5}` prevents narrowing conversions (compile error).

Idioms: prefer brace initialization `{}`; use `const` when the value won't change; use `auto` when the type is obvious or complex.

2. Operators

Precedence Table

Prec	Operator	Description	Assoc
1	<code>::</code>	scope resolution	L-R
2	<code>() [] . -> ++ --</code>	call, subscript, member access, postfix inc/dec	L-R
3	<code>++ -- + - ! ~ * &</code>	prefix inc/dec, unary, logical/bitwise NOT, deref, address-of	R-L
5	<code>* / %</code>	multiply, divide, modulus	L-R
6	<code>+ -</code>	add, subtract	L-R
7	<code><< >></code>	bitwise shift	L-R
8	<code>< <= > >=</code>	relational	L-R
9	<code>== !=</code>	equality	L-R
10	<code>&</code>	bitwise AND	L-R
11	<code>^</code>	bitwise XOR	L-R
12	<code> </code>	bitwise OR	L-R
13	<code>&&</code>	logical AND (short-circuits)	L-R
14	<code> </code>	logical OR (short-circuits)	L-R
15	<code>?: = += -= *= /= %=</code>	ternary, assignment, compound assignment	R-L
16	<code>,</code>	comma	L-R

Quick Reference

Operator	Example	Notes
<code>%</code>	<code>10 % 3 → 1</code>	integers only; remainder after division

Operator	Example	Notes
<code>++x / x++</code>	prefix returns new value; postfix returns old	prefer prefix in loops
<code>+= -= *= /= %=</code>	<code>x += 5</code> \equiv <code>x = x + 5</code>	compound assignment
<code>?:</code>	<code>cond ? a : b</code>	ternary; avoid nesting
<code>sizeof(T)</code>	<code>sizeof(int)</code> \rightarrow 4 (typical)	size in bytes; compile-time
<code>static_cast<T>(x)</code>	<code>static_cast<double>(n)</code>	safe C++ cast (e.g. int to double)

Traps: `=` vs `==` in conditions (`if (x = 5)` assigns, always true). `!` binds tighter than `>`, so `!score > 80` is `(!score) > 80`. Bitwise operators have lower precedence than `==`: `x & 0xFF == 0` parses as `x & (0xFF == 0)`. Don't confuse `&&/||` (logical) with `&/|` (bitwise).

Idioms: use parentheses when precedence is not obvious; integer division truncates (`10/3` \rightarrow 3).

3. Formatting

`#include <print>` for `std::println / std::print`; `#include <format>` for `std::format`

Signature	Description
<code>std::println("{} ", val)</code>	print with newline
<code>std::print("{} ", val)</code>	print without newline
<code>std::format("{} ", val)</code>	returns <code>std::string</code>
<code>std::println(file, "{} ", val)</code>	print to an output file stream

Format Specification: `{index:fill align width .precision type}`

Spec	Example	Result
fixed 2 decimals	<code>{:.2f}</code>	3.14
right-justify, width 10	<code>{:>10}</code>hello
left-justify, width 10	<code>{:<10}</code>	hello.....
center, width 10	<code>{:^10}</code>	..hello...
fill with *, right, width 6	<code>{:*>6}</code>	***abc
explicit arg index	<code>{0:} {1:}</code>	args by index
dynamic width	<code>{:>{}}</code>	2nd arg = width
decimal (default)	<code>{:d}</code>	42
hexadecimal	<code>{:x} / {:#x}</code>	2a / 0x2a
octal	<code>{:o} / {:#o}</code>	52 / 052
binary	<code>{:b} / {:#b}</code>	101010 / 0b101010
always show sign	<code>{:+}</code>	+42 or -42
space for positive	<code>{: }</code>	.42 or -42

Traps: literal braces require doubling: `format("{x}")` produces `{x}`. `std::cout` prints `bool` as 0/1; `std::println` prints `true/false`. Too few format args is a compile error; too many silently ignores extras. Displayed values are rounded, not truncated.

Idioms: prefer `std::format` over `+` for building strings; prefer `std::println` over `std::cout`.

4. Functions

```
int sum(int a, int b);           // returns int by value
void foo(int x);                 // by value (copies)
void foo(const std::string & s); // by const ref (no copy, read-only)
void foo(std::string & s);       // by ref (modifiable)
```

```
double avg(const std::vector<double> & v); // returns double
std::string greet(std::string_view name); // returns string
```

Pattern	When to Use
pass by value	fundamental types (<code>int</code> , <code>double</code> , <code>char</code> , <code>bool</code>), <code>string_view</code>
pass by <code>const &</code>	large objects you don't need to modify (<code>string</code> , <code>vector</code>)
pass by <code>&</code>	objects you need to modify, streams
return by value	the default; compiler optimizes away copies
return <code>void</code>	function has no return value

Traps: default parameters must be at the end; once you default one, all that follow must also be defaulted. A function in a header without `inline` breaks ODR if included in multiple files.

Idioms: pass small types by value, large types by `const &`; if the function needs its own copy, take by value to get an automatic copy.

5. iostreams

`#include <iostream>` for `cin/cout/cerr`; `#include <fstream>` for file streams; `#include <sstream>` for string streams

Stream	Description
<code>std::cin</code>	standard input (keyboard)
<code>std::cout</code>	standard output (buffered)
<code>std::cerr</code>	standard error (unbuffered)

Signature	Description
<code>std::cin >> var</code>	extract from standard input
<code>std::cout << val << '\n'</code>	insert to standard output
<code>std::getline(stream, str)</code>	read entire line into <code>string</code>
<code>std::getline(stream, str, delim)</code>	read until <code>delim</code> into <code>string</code>
<code>if (stream)</code>	check stream is in good state
<code>stream.clear()</code>	clear error flags
<code>stream.ignore(max, '\n')</code>	discard remaining input (see below)
<code>std::ifstream file{name}</code>	open for reading
<code>std::ofstream file{name}</code>	open for writing (truncates!)
<code>std::ofstream file{name, ios::app}</code>	open for appending
<code>std::istringstream iss{str}</code>	input from a string (useful for testing)
<code>std::ostringstream oss{}</code>	build a string via <code><<</code> , extract with <code>oss.str()</code>

File Mode	Description
<code>std::ios::in</code>	input (default for <code>ifstream</code>)
<code>std::ios::out</code>	output (default for <code>ofstream</code>)
<code>std::ios::app</code>	append
<code>std::ios::binary</code>	binary mode

Full ignore call: `stream.ignore(std::numeric_limits<std::streamsize>::max(), '\n')` (needs `<limits>`)

Traps: `>>` leaves the newline in the buffer; a following `getline` reads an empty string. Fix: call `ignore()` between `>>` and `getline`, or use `std::getline(std::cin >> std::ws, str)`. `while (!stream.eof())` is almost always wrong – `eof()` is only set after a failed read, causing the last item to be processed twice; use `while (stream >>`

var) or while (std::getline(stream, line)) instead. Non-numeric input puts cin in a fail state; must clear() then ignore() to recover. Default ofstream truncates the file. Always check if (file) after opening.

Idioms: write functions taking std::istream & / std::ostream & so they work with cin/cout, files, and stringstream. Files close automatically via RAII when the stream goes out of scope.

6. Strings and string_view

```
#include <string>; #include <string_view>
```

std::string

Signature	Description
std::string s{"abc"}	construct from literal
std::string s(3, 'A')	"AAA" – n copies of a char
s.find(ch)	size_t position, or npos if not found
std::string::npos	size_t max value – “not found” sentinel
s.substr(pos, count)	returns string (copies)
s.size() / s.length()	size_t number of characters
s.empty()	bool – true if empty
s.front() / s.back()	char & – first / last character
s.push_back(ch)	append a character
s[i]	access by index (no bounds check)
s + other	concatenation (creates temporaries)
s += str / s.append(str)	append in place (no temporaries)
s.insert(pos, str)	insert string at position
s.erase(pos, count)	remove count chars starting at pos
s.replace(pos, count, str)	replace substring
s.c_str()	const char * for C interop
s.starts_with(x)	bool (C++20)
s.ends_with(x)	bool (C++20)
s.contains(x)	bool (C++23)

Conversions

Signature	Description
std::stoi(str)	int from string (throws on failure)
std::stol(str)	long from string
std::stoll(str)	long long from string
std::stoul(str)	unsigned long from string
std::stoull(str)	unsigned long long from string
std::stof(str)	float from string
std::stod(str)	double from string
std::stold(str)	long double from string
std::to_string(num)	string from numeric value

std::string_view

Signature	Description
std::string_view sv{str}	non-owning view of a string
sv.substr(pos, count)	returns string_view (no copy)

Traps: all sto* functions throw std::invalid_argument if the string is not a number and std::out_of_range if the value is too large. Modifying the underlying string invalidates any string_view of it (dangling, undefined behavior). String + concatenation creates temporaries; prefer std::format.

Idioms: use `std::string_view` by value for read-only string parameters; use `const std::string &` when you need a guaranteed `std::string`. Use "hello"s suffix (with `using namespace std::string_literals`) to get a `std::string` from a literal.

7. Arrays and Vectors

```
#include <array>; #include <vector>
```

`std::array` – fixed size

Signature	Description
<code>std::array<double, 5> a{}</code>	5 doubles, all 0.0
<code>std::array a{1, 2, 3}</code>	CTAD deduces type and size
<code>a[i] / a.at(i)</code>	T & – access (at throws on bad index)
<code>a.size()</code>	<code>size_t</code> number of elements
<code>a.begin() / a.end()</code>	iterator to first / past-end

`std::vector` – dynamic size

Signature	Description
<code>std::vector<int> v{}</code>	empty vector
<code>std::vector v{1, 2, 3}</code>	CTAD
<code>std::vector<int> v(10, 0)</code>	10 elements, all 0
<code>v.push_back(val)</code>	add to end (may reallocate)
<code>v.emplace_back(args...)</code>	construct in place at end
<code>v.insert(it, val)</code>	insert at position
<code>v.erase(it)</code>	erase single element
<code>v.erase(first, last)</code>	erase range
<code>v[i] / v.at(i)</code>	T & – access (at throws on bad index)
<code>v.size()</code>	<code>size_t</code> number of elements
<code>v.empty()</code>	<code>bool</code> – true if no elements
<code>v.front() / v.back()</code>	T & – first / last element
<code>v.clear()</code>	remove all elements (size becomes 0)
<code>v.resize(n)</code>	change element count (adds default or removes)
<code>v.pop_back()</code>	remove last element
<code>v.reserve(n)</code>	pre-allocate space (no new elements)
<code>v.capacity()</code>	<code>size_t</code> current allocated capacity
<code>v.data()</code>	T * pointer to underlying array (C interop)

`std::span` – non-owning view of contiguous data (C++20)

```
#include <span>
```

Signature	Description
<code>std::span<int> s{arr}</code>	view of array or vector
<code>std::span<int, 5> s{arr}</code>	fixed-extent (size known at compile time)
<code>s.size() / s[i]</code>	<code>size_t</code> count / element access
<code>s.subspan(offset, count)</code>	<code>span</code> – sub-view

Range-based for:

```
for (const auto & elem : container) { ... } // read-only
for (auto & elem : container) { ... } // modifiable
```

Traps: `vector<int>{2, 5}` has two elements (2 and 5); `vector<int>(2, 5)` has two elements both equal to 5 – braces vs parentheses matter. `push_back` may reallocate and invalidate all iterators. `[]` does not check bounds.

Idioms: use `vector` as the default container; use `array` when size is known at compile time; call `reserve()` before many `push_back` calls.

8. Algorithms

`#include <algorithm>`; `#include <numeric>` for `accumulate`

Range Algorithms (C++20, `std::ranges::`)

Signature	Description
<code>ranges::sort(cont)</code>	<code>void</code> – sort in place
<code>ranges::sort(cont, comp)</code>	sort with comparator (e.g. <code>greater<>{}</code>)
<code>ranges::sort(cont, {}, &T::mem)</code>	sort by member (projection)
<code>ranges::find(cont, val)</code>	iterator to first equal element, or <code>end()</code>
<code>ranges::find_if(cont, pred)</code>	iterator to first match, or <code>end()</code>
<code>ranges::count_if(cont, pred)</code>	<code>size_t</code> count of matching elements
<code>ranges::minmax(cont)</code>	<code>{.min, .max}</code> struct
<code>ranges::any_of(cont, pred)</code>	<code>bool</code> – true if any element matches
<code>ranges::all_of(cont, pred)</code>	<code>bool</code> – true if all match
<code>ranges::none_of(cont, pred)</code>	<code>bool</code> – true if none match
<code>ranges::for_each(cont, fn)</code>	apply <code>fn</code> to each element
<code>ranges::transform(in, out, fn)</code>	apply <code>fn</code> , write to <code>out</code> iterator
<code>ranges::copy(cont, out)</code>	copy elements to <code>out</code> iterator
<code>ranges::copy_if(cont, out, pred)</code>	copy matching elements
<code>ranges::reverse(cont)</code>	<code>void</code> – reverse in place
<code>ranges::max_element(cont)</code>	iterator to largest element
<code>ranges::min_element(cont)</code>	iterator to smallest element
<code>ranges::generate(begin, end, fn)</code>	<code>void</code> – fill range from a callable
<code>ranges::generate_n(inserter, n, fn)</code>	iterator past last generated
<code>ranges::shuffle(cont, gen)</code>	<code>void</code> – random reorder using engine

`std::back_inserter(container) (<iterator>)` – output iterator that calls `push_back`; used with `copy`, `copy_if`, `transform`, `generate_n`.

Classic Algorithms

Signature	Description
<code>std::count_if(begin, end, pred)</code>	<code>size_t</code> count of matching
<code>std::accumulate(begin, end, init)</code>	<code>T</code> sum of elements (<code><numeric></code>)
<code>std::iota(begin, end, start)</code>	fill with <code>start</code> , <code>start+1</code> , ... (<code><numeric></code>)
<code>std::remove_if(begin, end, pred)</code>	iterator to new logical end

Erase (C++20)

Signature	Description
<code>std::erase_if(container, pred)</code>	<code>size_t</code> count of erased elements

Common Predicates and Comparators

`#include <functional>` for comparators; `#include <cctype>` for character classification

Predicate	Description
<code>std::less<>{}</code>	<code>a < b</code> (default for <code>sort</code>)
<code>std::greater<>{}</code>	<code>a > b</code> (sort descending)
<code>std::equal_to<>{}</code>	<code>a == b</code>
<code>std::isalpha(ch)</code>	<code>bool</code> – letter (<code><cctype></code>)
<code>std::isdigit(ch)</code>	<code>bool</code> – digit (<code><cctype></code>)
<code>std::isspace(ch)</code>	<code>bool</code> – whitespace (<code><cctype></code>)
<code>std::isupper(ch) / std::islower(ch)</code>	<code>bool</code> – case check (<code><cctype></code>)

`std::toupper(ch) / std::tolower(ch)` (`<cctype>`) convert case; commonly used with `ranges::transform`.

Traps: `std::ranges::minmax` on an empty range is undefined behavior. `std::remove_if` does NOT resize the container; must follow with `container.erase(new_end, container.end())`. `std::accumulate` is in `<numeric>`, not `<algorithm>`. `<cctype>` functions take `int`, not `char`; passing a negative `char` is undefined behavior – cast to unsigned `char` first.

Idioms: prefer `std::erase_if` (C++20) over the erase-remove idiom; prefer `std::ranges::` algorithms over iterator-pair versions; use `std::back_inserter(container)` with generate algorithms.

9. Lambdas

Syntax: `[captures](parameters) { body }`

Capture	Meaning
<code>[]</code>	nothing
<code>[x]</code>	<code>x</code> by value (copy)
<code>[&x]</code>	<code>x</code> by reference
<code>[x, &y]</code>	mix
<code>[=]</code>	all used variables by value
<code>[&]</code>	all used variables by reference
<code>[this]</code>	current object's <code>this</code> pointer
<code>[p = std::move(ptr)]</code>	init-capture: move into the lambda

Adding `mutable` after `()` allows modifying by-value captures (the copy, not the original).

Generic lambdas use `auto` parameters: `[](auto x, auto y) { return x + y; }`.

Trailing return type: `[](int x) -> double { return x / 2.0; }` – needed when return type is ambiguous.

`std::function<ReturnType(Params)>` (`<functional>`) stores any callable with a matching signature.

Traps: each lambda has a unique type; use `auto` to store one. Capturing by reference: if the referenced variable goes out of scope, the lambda has a dangling reference. `[=]` can accidentally capture unintended variables.

Idioms: capture explicitly by name rather than `[=]` or `[&]`. Use named functions for anything longer than a few lines. Use lambdas as algorithm predicates:

```
std::erase_if(v, [](double x) { return x < 0.0; });
```

10. Ranges

`#include <ranges>`

View	Description
<code>std::views::filter(range, pred)</code>	keep elements matching predicate
<code>std::views::transform(range, fn)</code>	apply fn to each element
<code>std::views::take(range, n)</code>	first n elements
<code>std::views::take_while(range, pred)</code>	take while predicate is true
<code>std::views::drop(range, n)</code>	skip first n elements
<code>std::views::drop_while(range, pred)</code>	skip while predicate is true
<code>std::views::reverse</code>	reverse the range
<code>std::views::iota(start)</code>	infinite sequence: <code>start, start+1, ...</code>
<code>std::views::iota(start, end)</code>	bounded sequence: <code>start</code> to <code>end-1</code>
<code>std::views::enumerate(range)</code>	{ <code>index, value</code> } pairs (C++23)

Pipe Syntax

```
auto result = data | std::views::filter(pred) | std::views::take(5);
```

Converting to a Container

```
auto vec = std::ranges::to<std::vector>(view);           // C++23
std::vector<int> vec{view.begin(), view.end()};         // fallback
```

Traps: views are lazy – predicates are not called until the view is iterated. Views do not own data; if the underlying data goes out of scope, the view is invalid.

Idioms: use pipe `|` to chain views for readability; use views to avoid copying large datasets.

11. Random Numbers

```
#include <random>
```

Three components: seed, engine, distribution.

Signature	Description
<code>random_device rd{}</code>	non-deterministic seed source
<code>default_random_engine gen(rd())</code>	engine seeded randomly
<code>default_random_engine gen{42}</code>	fixed seed (reproducible)
<code>uniform_int_distribution dist{1, 6}</code>	integers in <code>[1, 6]</code>
<code>uniform_real_distribution dist{0.0, 1.0}</code>	doubles in <code>[0.0, 1.0)</code>
<code>normal_distribution dist{mean, stddev}</code>	normal/bell curve
<code>bernoulli_distribution dist{0.5}</code>	bool with given probability
<code>int roll = dist(gen)</code>	generate a number

All types above are in `std::`.

Traps: same seed → same sequence. Without `std::random_device`, you get the same numbers every run. Normal distribution can produce negative values.

Idioms: use `std::random_device` for varying seeds; use a fixed seed for reproducible testing; separate generation from usage by passing a `std::function<double()>` or lambda.

12. Structures and Classes

struct (all public by default)

```
struct Point {
    double x{};
    double y{};
};
Point p{3.0, 4.0};           // aggregate initialization
Point q{.x = 3.0, .y = 4.0}; // designated initializers (C++20)
auto [x, y] = p;           // structured bindings (C++17)
```

class (all private by default)

```
class Stock {
    std::string name{};
    double price{};
public:
    Stock(const std::string & name, double price)
        : name(name), price(price) {} // member initializer list

    std::string get_name() const { return name; } // const member function
    void set_price(double p) { price = p; }
    ~Stock() = default;
};
```

Concept	Notes
<code>public / private / protected</code> member initializer list	access specifiers : <code>name(n)</code> , <code>price(p)</code> {} – preferred over body assignment
<code>const</code> member function	can be called on <code>const</code> objects; promises not to modify state
<code>explicit</code> constructor	prevents implicit conversions from single-parameter constructors
aggregate initialization	works for structs with all public members and no constructors
<code>static</code> member	shared by all instances; accessed via <code>Class::member</code>
<code>friend</code> function	non-member that can access private members
<code>bool operator==(const T &) const = default</code> <code>auto operator<=>(const T &) const = default</code>	compiler-generated equality (C++20) generates all 6 comparisons (C++20)

enum class

```
enum class Color { red, green, blue };
Color c = Color::red;
```

Scoped, no implicit conversion to `int`. Use `static_cast<int>(c)` to convert.

operator« overload

```
friend std::ostream & operator<<(std::ostream & os, const Point & p) {
    return os << "(" << p.x << ", " << p.y << ")";
}
```

Define as `friend` inside the class to access private members.

Traps: providing any constructor removes the implicit default constructor. Calling a non-`const` member function on a `const` object is a compile error. Assigning in the constructor body initializes members twice. Unscoped `enum` values leak into the enclosing scope.

Idioms: prefer member initializer lists; mark getters `const`; mark single-parameter constructors `explicit`; use `.` for direct access, `->` through a pointer. Prefer `enum class` over `enum`. Use `= default` for `operator==` when memberwise comparison is correct.

13. `unique_ptr` and Move

```
#include <memory>; #include <utility> for std::move
```

```
std::unique_ptr
```

Signature	Description
<code>auto p = std::make_unique<T>(args...)</code>	create on heap
<code>p->member</code>	access member
<code>*p</code>	dereference
<code>p.get()</code>	<code>T *</code> – raw pointer without releasing ownership
<code>if (p)</code>	check not nullptr
<code>p.reset()</code>	release and set to nullptr
<code>p.release()</code>	<code>T *</code> – relinquish ownership (caller must delete)
<code>auto q = std::move(p)</code>	transfer ownership (<code>p</code> becomes nullptr)

Cannot be copied. Only one `unique_ptr` owns the object at a time.

Move Semantics

- `std::move(obj)` casts to an rvalue reference, transferring ownership
- moved-from object is in a valid but unspecified state

Traps: `std::move` on a `const` object does nothing. `return std::move(local);` pessimizes by preventing NRVO – just write `return local;`. A class containing a `unique_ptr` member cannot be copied. `release()` gives you a raw pointer you must manually `delete`.

Idioms: always use `std::make_unique` – never raw `new/delete`.

14. Exceptions and Expectations

```
#include <exception>; #include <stdexcept>; #include <expected> for std::expected
```

Throwing and Catching

```
try {
    throw std::invalid_argument("negative value");
}
catch (const std::invalid_argument & ex) {
    std::println("{} ", ex.what());
}
catch (const std::exception & ex) {
    std::println("error: {} ", ex.what());
}
```

Common Exception Types

Type	Header	Typical Use
<code>exception</code>	<code><exception></code>	base type; generic catch-all
<code>invalid_argument</code>	<code><stdexcept></code>	bad parameter values
<code>runtime_error</code>	<code><stdexcept></code>	errors detected at runtime (e.g. file I/O)
<code>out_of_range</code>	<code><stdexcept></code>	index/value out of bounds

Type	Header	Typical Use
<code>bad_expected_access</code>	<code><expected></code>	<code>.value()</code> on error <code>expected</code>
<code>bad_optional_access</code>	<code><optional></code>	<code>.value()</code> on empty <code>optional</code>

All types above are in `std::`.

Custom exceptions: derive from `std::exception` and override `what()` if needed.

noexcept

Mark functions that will never throw: `void f() noexcept;`. Enables compiler optimizations and signals intent to callers.

std::expected (C++23)

Returns a value or an error without exceptions.

Signature	Description
<code>std::expected<T, E></code>	holds <code>T</code> on success or <code>E</code> on error
<code>return value;</code>	return a success value
<code>return std::unexpected{err};</code>	return an error
<code>result.has_value() / if (result)</code>	<code>bool</code> – check for success
<code>result.value()</code>	<code>T</code> – access value (throws if error)
<code>result.error()</code>	<code>E</code> – access error

Traps: catch handlers are tried in order – put specific types (`invalid_argument`) before general types (`exception`), or the specific handler is never reached. An unhandled exception terminates the program. Calling `.value()` on an error `expected` throws `bad_expected_access`. You can throw any type (even `int`), but doing so is unconventional – always throw types derived from `std::exception`.

Idioms: always catch by `const &`. Wrap the body of `main` in `try/catch` to prevent exceptions from leaking to the user. Don't use exceptions for flow control. Reserve exceptions for truly exceptional conditions – for expected errors (e.g. bad user input), prefer `std::expected` or return a `bool`. Check conditions before throwing when possible (e.g. `std::filesystem::exists()` before opening a file).